Substructural Content

by

Daniel Kaplan

B.A. Philosophy & Mathematics, Franklin & Marshall College, 2012

Submitted to the Graduate Faculty of

the Department of Philosophy in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

University of Pittsburgh

2021

UNIVERSITY OF PITTSBURGH

DEPARTMENT OF PHILOSOPHY

This dissertation was presented

by

Daniel Kaplan

It was defended on

December 9, 2021

and approved by

Robert B. Brandom, Distinguished Professor of Philosophy

James Shaw, Associate Professor of Philosophy and Department Chair

Erica Shumener, Assistant Professor of Philosophy

Greg Restall, Shelby Cullom Davis Professor of Philosophy, University of St. Andrews

Substructural Content

Daniel Kaplan, PhD

University of Pittsburgh, 2021

It is nearly ubiquitous for philosophers interested in meaning and consequence to first provide a semantics for sentences and next define what it means for sentences to follow from one another. Structural features of consequence such as monotonicity, transitivity, contraction, and reflexivity are rarely acknowledged except insofar as they are presupposed by whatever we intend with “follows from”. Since “follows from” succeeds sentence meaning, structural features are taken to presuppose robust constraints on the latter. My dissertation argues that this setup is mistaken; it misunderstands the connection between sentence meaning and what structural features require of content. For example, monotonicity is taken to follow from strong assumptions concerning the compositionality of sentence meaning. As a result the setup assumes that substructural consequence relations require different understandings of content. If we reject this setup and posit a closer relationship between meaning and consequence, then we are better able to understand what sorts of constraints are placed on semantic content by structural rules. I develop a view in which the two are maximally close: meaning just is contribution to consequence. I argue that this notion of content, when precisified yields a completely tractable, robust semantics. Refiguring the relationship between meaning and consequence also yields surprising insights for the philosophy of language and meta-ethics.

Keywords: philosophy of language, philosophical logic, substructural logic, inferentialism, defeat, defeasible reasoning, radical contextualism, semantic minimalism, moral particularism.

Table of Contents

List of Tables

List of Figures

Preface

Writing a dissertation is a lot like building a house. If the house were made of words, and instead of tools you had thoughts and paper, and no one was going to live there, it was just to be used for a single dinner party with a small number of guests. Also the guests alternate between the highest of praise and harshest critiques of the menu you've chosen to prepare and how you've chosen to prepare it. Of course a resourceful homebuilder can find a way to sell the house rather than leave it vacant or at least repurpose the materials to build more houses later on. There isn't really a resale market for dissertations though. Maybe it's not like building a house.

There are a number of people[1] I should thank and acknowledge. First, thanks to my committee—Bob Brandom, James Shaw, Erica Shumener, and Greg Restall—whose patience has been boundless during this process.

In particular, I should thank my adviser, Bob Brandom, weekly meetings with whom together with feedback on countless pages I sent him made the formulation and completion of this project possible. I recall once after a colloquium talk going to dinner with the speaker, some grad students, and some professors (among them James). As is typical, the speaker asked the grad students present, "what do you work on?", "who's on your committee?", and so forth. Eventually the speaker asked me (or the topic somehow turned towards), what it was like to work with Bob—to which I probably gave a generic answer. James volunteered that Bob is known for meeting with his students weekly, requiring of them something written each week,[2] and offering often quite attentive and detailed feedback on that writing. What stuck out to me about this interaction, was the bafflement bordering on indignation that the speaker seemed to show upon learning this. How could Bob do such a thing? Supposing James wasn't speaking exaggeration bordering on untruth, it set a standard no one could hope to achieve. The point of this anecdote is to emphasize the care and attention Bob gives to his students, and the large degree to which I benefited from this. I hope that the

[1] With a small exception towards the end, I avoid thanking things. I recently calculated that I drank enough coffee during my graduate students to *easily* fill a six-person hot-tub with coffee.

[2] A requirement that I must report that I flouted far too often.

document before you is proof of that.

During the research and writing of the dissertation I regularly participated in a research group ROLE (Research on Logical Expressivism) organized by Bob Brandom. Some of the work that appears here was presented there. I am grateful for valuable feedback I received over the years, in particular to: Bob Brandom, Ulf Hlobil, Rea Golan, Ryan Simonelli, Shuhei Shimamura, and Patrick Chandler.

Some of the material in the dissertation was presented to audiences at a number of conferences and discussed with fellow presenters and audience members. I am forgetting many names, but I do recall having stimulating conversations on more than one occasion with Dave Ripley, Ulf Hlobil, Jarda Peregrin, Lucas Rosenblatt, and Damian Szmuc. I should also thank John Horty: a conversation with whom led to the section on particularism (§4.2)—though I don't think that was at all his intention—and Alex Kocurek, who helped tremendously with getting the formulation of the precisification of logical expressivism right (§2.2). In addition, I am grateful for discussion and feedback at the following venues: the "Defeasible Inference in Philosophy and Artificial Intelligence" Conference (UCLA 2019); Logica 2017 and 2018 (Hejnice, Czech Republic);[3] "PhDs in Logix X" Conference (Prague, 2018); the 5th and 7th Workshop on Philosophical Logic of the Argentinean Society of Philosophical Analysis (SADAF) (Buenos Aires, 2016 and 2018); and the "Logics of Consequence" Workshop (Concordia, Montreal, 2017).

During a visit to the University of Leipzig in 2018, I benefited from conversation with Pirmin Stekeler-Weithofer, as well as from the opportunity to teach a graduate seminar on inferentialism and philosophical logic. Both of these helped clarify a very early conception of the project.

Partial chapter drafts were presented to the dissertation seminar at the University of Pittsburgh. I am grateful for the feedback and ideas generated by my fellow graduate students during those sessions. Particularity risks omission, but I recall particularly helpful comments from: Tom Breed, Max Tegtmeyer, Alnica Visser, Josh Fry, Alessandra Buccella, Stephen Mackereth, Travis McKenna, and Dan Webber. James Shaw and Jed Lewinsohn, who each ran the seminar at different points while I attended, gave additional helpful com-

[3]Some of the material in §2.2 appeared in Kaplan (2018), for which I retained the copyright.

ments on several drafts in connection with this.

Stephen Mackereth and Lucas Rosenblatt each provided extremely detailed and useful feedback on drafts of the material on contraction in the dissertation (§3.2). Comments from James Shaw really helped to clarify the dialectic I was grasping for (which I hope makes the current presentation much clearer than it was in earlier drafts). Without their help—for which I am extremely grateful—I'm not sure that section would have turned out as nicely as (I take) it has.

I don't know if I would have been able to make it through this process without the friendship and stimulating conversation—on all topics philosophical or not—from my fellow graduate student, Max Tegtmeyer.

Finally, to my wife, à mon amour, Ariel Kornhauser: thanks for nothing.[4]

In addition to all of this expression of gratitude, I'd like to dedicate this dissertation to Satan: desecrated be his name.[5]

[4] She asked me to put this in here.

[5] I should clarify that Satan is the name of one of my cats. Likewise, I should also dedicate the dissertation to MayoChup and Frisky Lincoln, the last of whom didn't make it to the end (RIP). See Figure 1.

Satan

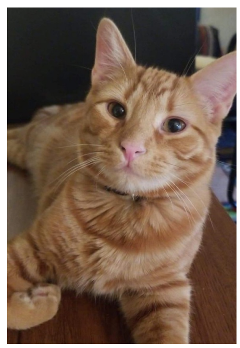

MayoChup

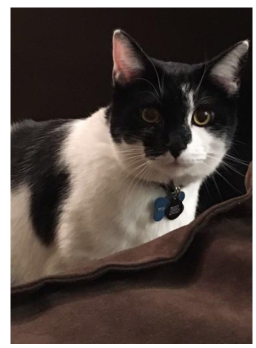

Frisky Lincoln (RIP)

Figure 1: A few good men.

1.0 Introduction

The goal of this dissertation is to put forward a theory of content and see how far it can take us. The theory of content amounts to the suggestion that we understand the content of a sentence in terms of its contribution to good implication. The twist is that I also acknowledge that good implications may be radically substructural, where "structural" refers to structural proof rules in sequent calculi (e.g. such rules as monotonicity, contraction, transitivity, and reflexivity).[1] Substructural then means capable of violating these rules. I therefore call such content *substructural content*, and I suggest that this theory can take us quite far.

1.1 An Outline

The dissertation is divided into three parts. In the first part (Ch. 2), I articulate my theory of meaning and explain the relevant background. The main goal of this chapter, aside from putting the theory on the table, is to convince the reader that the formal semantics I construct alongside the theory is completely tractable. I do this by showing that we can introduce the standard logic connectives and that simple and familiar proof rules (in the form of a sequent calculus) are sound and complete for the semantics. I also prove some results for the proof system that will be of particular interest to philosophers and logicians interested in inferentialism. Using the proof system, I give a precisification of logical expressivism. Logical

[1] See Figure 2 (p. 11, reprinted in this footnote):

Monotonicity

$$\frac{\Gamma \vdash \Theta}{\Delta, \Gamma \vdash \Theta}\text{ L-MO} \qquad \frac{\Gamma \vdash \Theta}{\Gamma \vdash \Theta, \Lambda}\text{ R-MO}$$

Transitivity/Cut and **Reflexivity**

$$\frac{\Gamma \vdash \Theta, A \qquad A, \Gamma \vdash \Theta}{\Gamma \vdash \Theta}\text{ Shared-Cut} \qquad \frac{}{\Gamma, A \vdash A, \Theta}\text{ Reflexivity}$$

Contraction

$$\frac{A, A, \Gamma \vdash \Theta}{A, \Gamma \vdash \Theta}\text{ L-Contr.} \qquad \frac{\Gamma \vdash \Theta, A, A}{\Gamma \vdash \Theta, A}\text{ R-Contr.}$$

expressivism is the view that we should understand the meaning of logical connectives in terms of their expressive role. In light of the precisification, I show that the connectives introduced in my formal semantics are expressive in exactly this sense: they allow us to express *exactly* what follows from what. Following this I also show a second sense in which the machinery can be said to be expressive: we are able to mark in the object language where structural rules hold. This result not only advances the thesis of logical expressivism, but it also helps to show the tractability of the formal project. I allow that implication be non-monotonic, for example, but we are able to express where monotonicity *does* hold by the use of a sentential operator ($\boxed{M}$). One particularly interesting feature that I show can be marked in this way is captured by the natural language locutions "literally" or "strictly speaking". I show that we may introduce an operator into the object language ($\boxed{L}$) which express when implications are *fully* structural (i.e. monotonic, transitive, contractive, and reflexive). My claim is that the notion of content recaptured by this connective corresponds to more traditional understandings of content. $\boxed{L}p$ is the "minimal proposition" expressed by p.[2]

Finally, I examine attempts to account for the non-monotonicity of implication in terms of defeasible reasoning (and thus in terms of defeaters). I argue that we simply cannot account for the non-monotonicity of reasoning in these terms because we are unable to produce an account of defeaters or the phenomenon of defeat. Thus, I advocate that we should think of this phenomenon in terms of how a set of considerations hangs together; that is, in terms of the non-monotonicity of reasoning. This is meant to form a partial argument in favor of going substructural.

Having shown that we can construct a tractable formal semantics of meaning in terms of contribution to good implication and having shown why we ought to understand this phenomenon in terms of its substructurality, I next articulate what is distinctive of this approach (Ch. 3). By understanding meaning in terms of contribution to good implication, I argue that we must *push the substructurality all the way down into the content*. I argue that many other substructural logics fail to do this and thus do not take the substructurality

[2]This result will be particularly interesting when considering the third part (Ch. 4) of the dissertation, where I discuss literal meaning and minimal propositions.

of implication seriously. They are committed to what I call the *Assumption of Structurality*, which is the idea that in order to account for violations of structural rules, we must postulate some further layer of content which stands in fully structural relations of implication. It is through this postulation that a sentence p may be involved in a consequence relation which is non-monotonic (to continue the example):

$$\Gamma \vdash p \quad \text{but} \quad \Gamma, \Delta \nvdash p,$$

because there is some further layer of content that p expresses. In the context of Γ, p expresses p_1 and in the context of $\Gamma \cup \Delta$, it expresses p_2. Then:

$$\Gamma \vdash p_1$$
$$\Gamma, \Delta \nvdash p_2,$$

is not an example of a violation of monotonicity.[3] If we understand the content of a sentence in terms of its contribution to good implication such that that contribution is pushed all the way down into the content, such a postulation is not needed to understand how a sentence can be involved in implications which are substructural. Along the way, I prove some interesting results concerning substructural approaches to paradoxes. In particular, I provide some insight into the relationship between non-contractive, non-transitive, and non-reflexive solutions to paradox.

Finally, in the third part (Ch. 4), having shown that we can construct a tractable formal semantics and having articulated what is distinctive of pushing the substructurality of implication all the way down into the content, I show that this idea can be used to open up logical space in diverse regions of philosophy. I look at the semantic minimalism and radical contextualism debate in the philosophy of language, and the moral particularism and moral generalism debate in meta-ethics.

The former debate concerns whether there is such a thing as the "literal meaning" of a sentence and whether that thing is *what is said* by a sentence when it is asserted. Against this contextualists argue that what is said by a sentence may vary (perhaps radically so)

[3]Provided that $\Gamma, \Delta \vdash p_1$ and $\Gamma \nvdash p_2$, but this is precisely what the postulation is intended to make room for.

from occasion to occasion. I argue that understanding the content of a sentence in terms of its contribution to good implication provides us a way of both acknowledging the complexity of behavior that a sentence may exhibit, while also allowing us to understand the sentence as expressing the same content from occasion to occasion. To do this I expose a common presupposition underlying the debate and show how denying it opens up new logical space in the debate.

In the latter debate, moral particularists and moral generalists disagree about whether normative verdicts concerning actions are governed by moral principles. A moral principle in this debate is a rule which connects features of an action—typically statable in non-normative vocabulry—with reasons for/against that action, and eventually (perhaps through further principles governing the combination of such reasons) issues normative verdicts for the action. Moral particularists argue that we cannot articulate such principles owing to the complexity of the ways in which various reasons and considerations may interact from occasion to occasion. I show that a common presupposition underlying the debate may be denied in order to open up new logical space. Part of what we must deny to get there is the idea that a consideration is only a reason if a principle is able to explain what it is a reason for or against. Instead, if we understand reasons as that with which we reason (considerations active in our reasoning), then we can embrace a substructural notion of principle and thus reason that allows us to understand how the complexity of normative verdicts may be undergirded by stable contents.

In closing, I show that there is a deep commonality between these two debates. The crude formulation of this is that moral particularism is just radical contextualism concerning normative judgments; or, that radical contextualism is just particularism concerning truth conditions. Part of what I offer is more nuance concerning this crude approximation and thus more insight.

Understanding meaning in terms of its contribution to good implication, where such implication may be radically substructural can be used to construct a tractable formal semantics. This tractable semantics understands the substructurality of implication to be itself part of the content of a sentence. Understanding the content of a sentence in this way, I hope to have shown can do some real work outside of semantics and the philosophy of logic.

I hope to show how it can do this in two disparate debates in philosophy. I believe that there is room for it to make contributions in a number of other areas.

2.0 Meaning as Contribution to Good Implication

This chapter puts forward a candidate for the meaning of a sentence. The goal is to see what this candidate has going for it and how far it can take us. In particular, I put forward an *inferentialist* theory of meaning for sentences. This is a theory of meaning according to which the content of a sentence is be identified with that sentence's role in reasoning; i.e., the role it plays as the premise and conclusion of arguments. The account I want to explore takes this identification quite seriously: we simply understand a sentence in terms of this role.

Further, I believe that the implications in which a sentence is involved are substructural: meaning they have the potential to violate rules of inference that are often taken to be constitutive (of inference or consequence). I have in mind such structural features as monotonicity: if B follows from A, then B follows from A together with an arbitrary additional sentence C. In symbols:

$$A \vdash B \Rightarrow A, C \vdash B.$$

I also have in mind transitivity: if C follows from B and B follows from A, then C must also follow from A.

$$(A \vdash B \text{ and } B \vdash C) \Rightarrow A \vdash C.$$

Tarski, for example, understands consequence in terms of at least these two notions.[1] When I say that I want to investigate a notion of implication that is substructural, what I mean is that sentences ought to be able to stand in relations of consequence that potentially violate these rules. And they ought to be able to violate these rules in potentially radical ways.

This chapter advances such a semantic project. I start by providing a little motivation concerning substructural logic and explaining what it means to develop a theory of meaning for sentences. Following this I introduce inferentialism and my formal semantics according to which the meaning of a sentence is to be understand in terms of its contribution to good implication. Because the theory of meaning I develop involves radical substructurality—

[1]The other two structural rules that are baked into Tarski's definition are contraction and reflexivity, but more on these later.

radical here means that violations of structural rules are not accounted for by some "deeper" feature of the content—it might seem that the theory of content I develop isn't particularly tractable. In order to show that the semantics is tractable, I:

- Show that a fairly simple sequent calculus is sound and complete with respect to it.
- Develop a representation theorem for reproducing theories (in particular logical theories) with said sequent calculus.
- Along the way, I show how the formal machinery developed can help make sense of a topic of importance for inferentialists: logical expressivism.

Finally, in the close of the chapter, I provide an argument for going substructural. In particular, I try to show that attempts to account for violations of monotonicity in terms of "defeasible reasoning" don't get things right. I do this by arguing that accounts of defeat (in e.g. epistemology) fail to successfully isolate the target notion. Instead, I argue, we should understand violations of monotonicity in terms of how various considerations hang together. But the way that considerations hang together is wildly unpredictable. If this is right, then an account which allows us to understand non-monotonic implication without appeal to some "deeper" features would be a good thing indeed.

2.0.1 Substructural Logic

This dissertation is concerned with substructural logics. I should therefore take some time to explain what the "structure" in substructural refers to and thus what the "sub" in substructural refers to. Structural logic is a branch of proof theory concerned with so-called "structural" rules of proof. Its main concern is with sequent calculi, but the questions motivating this branch of logic can already be seen in studies of natural deduction. For example, if we have a proof,

$$\cfrac{\cfrac{A}{\vdots}}{C}$$

which begins with the undischarged assumption A and ends with C, we typically say that on the basis of A, we can prove C (or A proves C, $A \vdash C$ for short). But this informal notion is rather vague. For example, colloquially we might think that "on the basis of A and B, we can prove C" (since we can prove it from A), hence $A, B \vdash C$.[2] We might therefore want to distinguish a sense in which $A \vdash C$ but $A, B \nvdash C$,[3] since B is irrelevant[4] to the proof.[5]

Likewise, we might have a proof:

$$\dfrac{\dfrac{A}{\vdots} \qquad \dfrac{A}{\vdots}}{C}$$

It seems that $A, A \vdash C$, and there's clearly a colloquial sense in which, "on the basis of A, we can prove C", but this colloquial sense fails to mark that we've relied on A in two places. Hence, we might want to be careful distinguish that $A, A \vdash C$ but $A \nvdash C$. This is important if A can be used in ambiguous or divergent ways and we might want to keep track of that.

Particularly clear examples of this can be found among the semantic paradoxes. Paradoxes of the general sort I have in mind tend to treat the same object in divergent ways and this treatment can be easily missed if we allow ourselves to ignore cases where this happens. For example, consider the liar paradox. For the purposes of the example all that is important is that we allow ourselves to substitute α with $\tau(\alpha)$ for any α (τ is our truth predicate) as well as the following sort of substitution (λ is the liar sentence):[6]

$$\lambda = \tau(\neg\lambda).$$

Now, suppose we reason:

$$\dfrac{\lambda \qquad \dfrac{\dfrac{\lambda}{\tau(\neg\lambda)}}{\neg\lambda}}{\bot}$$

[2]It is conventional to use commas to separate premises in sequent calculi. This is especially important in structural proof theory, since we want to be careful to distinguish the kinds of structure that a collection of premises has from the kind of structure that e.g. a *set* of premises has (since sets come with their own structural assumptions built in).

[3]The '$\nvdash$' is shorthand in the meta-theory for: "it is not the case that $\ldots \vdash$ ___".

[4]I don't intend to advance a project in relevance logic. I am using what I take to be intuitive examples to help make intelligible what it means to relax these structural rules.

[5]The same questions can be raised for individual connectives. I.e. can we introduce a conditional by "discharging" a non-existent assumption?

[6]I also am going to assume standard natural deduction rules for negation and '$\bot$':

$$(n)\ \dfrac{\dfrac{\dfrac{[A]}{\vdots}}{\bot}}{\neg A}\ \neg\mathrm{I}\ (n) \qquad\qquad \dfrac{A \qquad \neg A}{\bot}\ \bot\mathrm{I}$$

A question we might ask is whether $\lambda \vdash \bot$. In a sense this seems right, but we might want to be careful to note that λ is playing two roles in this proof. In the first case it stands for some sentence, in the other case, we have substituted it multiple times (we are relying on its definitional equivalence and not on its status as just some sentence like any other). This is common to paradoxical reasoning. We often end up treating the same thing as just a mere sentence (like any other) or as something thicker: what it is defined as. If we want to mark this, we might want to be careful to write $\lambda, \lambda \vdash \bot$ since λ's use is potentially equivocal. If we allow ourselves (in the natural deduction setting) to simultaneously discharge both instances of λ, then we can derive $\neg\lambda$ (a surprising result). If we allow a few other minor substitution principles (that we can substitute λ for its definiens even when embedded in a larger sentence and double-negation elimination), then we can derive the following result (which relies on two instances of dual discharge):

$$\dfrac{\dfrac{{\scriptstyle(1)}\,\dfrac{[\neg\lambda] \quad \dfrac{\dfrac{{\scriptstyle(1)}\,\dfrac{[\neg\lambda]}{\neg(\tau(\neg\lambda))}}{\neg(\neg\lambda)}}{\lambda}}{\bot}}{\neg\neg\lambda}\,{\scriptstyle\neg\mathrm{I}\ (1)} \qquad \dfrac{{\scriptstyle(2)}\,\dfrac{[\lambda] \quad \dfrac{{\scriptstyle(2)}\,\dfrac{[\lambda]}{\tau(\neg\lambda)}}{\neg\lambda}}{\bot}}{\neg\lambda}\,{\scriptstyle\neg\mathrm{I}\ (2)}}{\bot}$$

Here is the paradigmatic problematic proof. From the definition of the liar ***alone*** (i.e. with no further undischarged premises) *follows* absurdity (i.e. $\bot$). Traditional accounts might try to deny some principle we've appealed to (i.e. that we may have self-referential definitions, or a truth predicate, etc.), but clearly the liar means something. The paradox itself is proof that we know what the sentence means in virtue of the fact that we reason with it. The problematic move, then, might just be that we have invoked λ (and $\neg\lambda$) equivocally and then glossed over that equivocation by discharging both instances simultaneously—note that if we don't allow dual discharge, the proof is simply $\neg\lambda, \lambda \vdash \bot$ (that contradictions imply falsum). We might want to therefore be careful to keep track of how many times we invoke a premise in order to guard against invoking it equivocally.[7]

Finally, suppose we have a proof from A to C, and then from C together with B to D:

[7] Another way of putting the point is that there's a choice point here. Either λ is the sort of thing that the meaning of a sentence could be (in which case we should deny contraction), or it is not (in which case we can keep contraction, but we have to explain why it is that we think it means something; how is it different from gibberish?).

$$\begin{array}{c} A \\ \hline \vdots \\ \hline C \end{array} \qquad\qquad\qquad \begin{array}{c} B \quad C \\ \hline \vdots \\ \hline D \end{array}$$

Bracketing the previous issue, would it be okay to say that we have a proof from A and B to D ($A, B \vdash D$)? In particular, couldn't we just 'cut and paste' the proofs together like this:

$$\begin{array}{c} \begin{array}{cc} & \begin{array}{c} A \\ \hline \vdots \\ \hline \end{array} \\ B & C \end{array} \\ \hline \vdots \\ \hline D \end{array}$$

Again, we might think this is fine. Especially since we rarely encounter cases where we couldn't just prove D straightaway from A, B. But mathematicians might want to note that we took a detour through C, especially if our aim is the study of proofs, we might want to distinguish a "direct" proof, from a proof that must first establish some powerful result (C) and then use that result to derive the intended result. Perhaps that powerful result is itself controversial, or, again, we might simply want to mark such a difference.[8] Statistical and probabilistic reasoning often has this form (it may be that A makes B likely, B makes C likely, but A does not make C likely).

In general, however, what a denial of cut amounts to is the idea that a sentence might not *do* the same thing (or play the same role) when it figures as the conclusion of an argument and as the premise of an argument (insistence on "cut" amounts to denying this possibility). For example, suppose that $A =$ "b is a bird" and $C =$ "b can fly" and we allow $A \vdash C$ (that C follows from A). Now suppose that $B =$ "b lives in Antarctica" and $D =$ "b is a penguin". Clearly we should also allow that $B, C \vdash \neg D$ since if b can fly and is from Antarctica it can't

[8] We can similarly sneak in other structural constraints from this one via connective rules. For example, maybe we don't like saying that $A, B \vdash C$ if B doesn't figure in the proof, but we are fine with our conditional behaving this way. Then just let C be $B \supset A$ and now B seems to play no actual role in the following proof (I suppose we have a way to reason from A to D):

$$\begin{array}{c} \begin{array}{cc} & \begin{array}{c} A \\ \hline B \supset A \end{array} \\ B & \end{array} \\ \hline A \\ \hline \vdots \\ \hline D \end{array}$$

The problem is that while $A \vdash B \supset A$ and $B, B \supset A \vdash C$ are both fine, $A, B \vdash D$ conceals an important middle step and there is no way to get this proof without that middle step unless one allows irrelevant premises.

Monotonicity

$$\frac{\Gamma \vdash \Theta}{\Delta, \Gamma \vdash \Theta}\text{ L-MO} \qquad\qquad \frac{\Gamma \vdash \Theta}{\Gamma \vdash \Theta, \Lambda}\text{ R-MO}$$

Transitivity/Cut

$$\frac{\Gamma \vdash \Theta, A \qquad A, \Gamma \vdash \Theta}{\Gamma \vdash \Theta}\text{ Shared-Cut}$$

Contraction

$$\frac{A, A, \Gamma \vdash \Theta}{A, \Gamma \vdash \Theta}\text{ L-Contr.} \qquad\qquad \frac{\Gamma \vdash \Theta, A, A}{\Gamma \vdash \Theta, A}\text{ R-Contr.}$$

Figure 2: Structural Rules

be a penguin. But cutting and pasting these together we get the result that: $A, B \vdash \neg D$, which we might not want. That b is not a penguin hardly seems to follow from its being a bird from Antarctica (its being such a bird might in fact be the most salient assumption).

The above three examples concern what are called monotonicity, contraction, and transitivity (here: shared-cut) in structural logic. If we use '$\vdash$' to encode proof, then they become the conditions on '$\vdash$' depicted in Figure 2 (the horizontal line indicates that the sequent below the line holds in case the sequent above the line holds) Substructural logics seek to deny some or all of these or to investigate under what conditions they should be suppressed or what results when they do. My suggestion is to investigate what effects such structural rules have on our understanding of content, generally, and whether we can furnish theories of substructural content.

2.1 Semantics

To start it's worth saying a few remarks about "the meaning of a sentence". Lewis (1970) famously distinguished between two senses in which we can offer a theory of meaning:

> "I distinguish two topics: first, the description of possible languages or grammars as abstract semantic systems whereby symbols are associated with aspects of the world; and, second, the description of the psychological and sociological facts whereby a particular one of these abstract semantic systems is the one used by a person or population. Only confusion comes of mixing these two topics." (Lewis, 1970, p. 19)

The first topic is typically called a "theory of meaning", the second a "meaning theory".[9] Meaning theories describe a particular interpretation of a particular language according to which items in that language are assigned meanings in a way that preserves some structure of the language. A theory of meaning says what form a meaning-theory must take (or sets some general constraints on meaning-theories). That is: a theory of meaning tells us what the elements of a language are assigned to, and what the relevant structure is. In this dissertation I provide a theory of meaning.

I provide a theory of the meaning of *sentences* for several reasons.[10] First, I agree with Frege's context principle that the constituents of a sentence acquire their content through their appearance in sentences (emphasis mine):

> "Accordingly, any word for which we can find no corresponding mental picture appears to have no content. But we ought always to keep before our eyes a complete proposition. *Only in a proposition have the words really a meaning.* It may be that mental pictures float before us all the while, but these need not correspond to the logical elements in the judgment. It is enough if the proposition taken as a whole has a sense; it is this that confers on its parts also their content." (Frege, 1980, §60)

It is only in the context of a sentence that words acquire their content. Thus, if one can put forward an account of the content of a sentence, then the meanings of sub-sentential expressions ought to be downstream from this.[11]

[9] As far as I can tell these terms originated with Dummett (1975, 1976). Their use is fairly ubiquitous now.

[10] Not least of which is that getting the sentential right is difficult enough. I hope to develop the work in this dissertation for quantification and sub-sentential expressions.

[11] Though this oversimplifies the matter.

Next, we ought to take the content of a sentence to be connected in some way with the (or some privileged) characteristic *use* of that sentence. For example, one (or perhaps *the*) basic function of language is communication. The most basic communicative act is taken to be assertion. So whatever content we express by asserting a sentence is what we should understand the meaning of the sentence to be.

Here there emerges a large divide among theories of meaning. Some understood assertion in terms of something like a norm for truth.[12] In that case, we should understand the content of a sentence in terms of its potential for truth/falsehood. According to this view, then, the content of a sentence is to be identified with its truth conditions (i.e. those circumstances under which it is true).[13]

In contrast to this, however, we might instead understand assertion not in terms of norms for its proper performance, but in terms of what we accomplish through assertion. Lewis (1979); Brandom (1983) suggest understanding assertion in terms of making a move in a language game; in particular the undertaking of a commitment. If we have a rational conception of language and communication, then the commitment is a commitment to justifying one's claim if challenged. According to such a conception of language, an essential aspect of a sentence's meaning is how it figures in reasoning: the sorts of claims it is capable of supporting under challenge and what sorts of further claims could help support its assertion (on challenge). What sorts of claims a sentence is capable of supporting, we should understand in terms of what consequences the sentence has, and what sorts of claims are capable of supporting a sentence, we should understand as those sentences which have the original sentence as a consequence. Taking these two aspects of a sentence to be constitutive of the sentence's meaning gives us a particular theory of meaning: an inferentialist theory of meaning.

[12] I am oversimplifying the matter. In particular, I am limiting my attention to normative conceptions of assertion. Many understand assertion, for example, in terms of the expression of a belief (or some other mental state). Then the meaning of a sentence corresponds to the structure of that mental state. Or perhaps the sentence simply refers to such a mental state (if we individuate mental states functionally, such a proposal is not that far fetched). At any rate, this takes me rather far afield from where I hope to go.

[13] I should note that when it comes to normative accounts of assertion, understanding assertion in terms of *knowledge* is a much more popular proposal. See Williamson (1996), who popularized this proposal. That assertion should be understood in terms of a norm for *truth* is put forward by MacFarlane (2011).

2.1.1 Inferentialism

Inferentialism is the view that the role that something (paradigmatically, a sentence) plays in reasoning, e.g. the role it serves as a premise and conclusion of arguments, is an essential element of that thing's meaning. In the central case, the meaning of a sentence is given by its role as a premise and a conclusion in argument, or as I shall say: the meaning of p just is the contribution that p makes to the goodness of implication. Thus, we might understand the meaning of p as specified in:

$$\begin{array}{ccc} p, \Gamma_1 \vdash \Theta_1 & \qquad\qquad & \Gamma_1 \vdash \Theta_1, p \\ p, \Gamma_2 \vdash \Theta_2 & & \Gamma_2 \vdash \Theta_2, p \\ \vdots & & \vdots \\ p, \Gamma_n \vdash \Theta_n & & \Gamma_n \vdash \Theta_n, p \\ \vdots & & \vdots \end{array}$$

On an inferentialist understanding of meaning, we treat the meaning of p as the contribution it makes to good inference above. This immediately introduces two constraints.

(Constraint One) The first constraint is that a sentence is only meaningful if it has a role as a premise and as a conclusion, (i.e. we must specify *two* lists). It's important that we specify two roles for at least two reasons. First, two sentences may play more-or-less the same role as a premise (or as a conclusion) but play different roles as conclusions. Second, the idea that a sentence might appear as a conclusion but never as a premise (or vice-versa) is unintelligible if we understand what sentences express to be rationally related to other sentences. We can appreciate both of these points clearly by considering two concepts that are extensionally equivalent (but differ intensionally: or two sentences with equivalent truth conditions, which express different senses, if one prefers to keep things sentential).

I introduce the following shorthand to talk about the contribution that, e.g. a sentence, makes to good implication. Since we'll need a way of noting the role that a sentence plays as a premise and as a conclusion, we'll have to note which one we're after:

$$\langle \{p\}, \emptyset \rangle^{\curlyvee} =_{df.} \{\langle \Gamma, \Theta \rangle | p, \Gamma \vdash \Theta\}, \qquad (p \text{ as premise})$$

$$\langle \emptyset, \{p\} \rangle^{\curlyvee} =_{df.} \{\langle \Gamma, \Theta \rangle | \Gamma \vdash \Theta, p\}. \qquad (p \text{ as conclusion})$$

which specify the contribution that p makes as a premise and conclusion, respectively, to the goodness of implication. Putting it all together then, I use double-brackets, '$[\![\cdot]\!]$" to denote the contribution of p in total:

$$[\![p]\!] =_{df.}= \langle \langle \{p\}, \emptyset \rangle^{\curlyvee}, \langle \emptyset, \{p\} \rangle^{\curlyvee} \rangle.$$

We might think of this as shorthand for the contribution to good implication that p made in the lists above. This constraint is equivalent in the valuational setting to the idea that a sentence needs specifiable truth-conditions. According to that account you can say many things about a sentence or how it might be spoken truly, but if you haven't specified truth conditions, you haven't specified its meaning.

The second constraint (below) also has a valuational correlate. This is the constraint that if two sentences have the same truth conditions then they are equivalent: that is, we must be extensional concerning truth conditions. To be precise: sentences are equivalent if they are true on exactly the same models. This constraint will end up imposing a kind of extensionality concerning meaning. This is especially important because we don't want to allow one sentence to have multiple *contents*; nor to allow *distinct* sentences the same content. The second constraint therefore fixes sentence individuation at the level of contribution to good implication.

(Constraint Two) A contribution to good implication fixes an equivalence class among inferential roles. Above I said that we had to distinguish two "roles" that a sentence plays in implication (two separate contributions it makes), but what counts as a role? I gave two lists, but is that what it means to be a role? Are there any constraints on what counts as contributing to the goodness of implication as either a premise or conclusion? Because I'm interested in relaxing structural rules of implication (in particular, contraction), I intend to be fairly liberal as regards what could count as a role in good implication. Even if I allow myself maximal liberality, however, one constraint still arises: contributions are equivalence classes of implicational roles. The intent of this is to fix equivalence of meaning at a certain level; namely the ground level: in terms of contribution to good implication. p (as a premise) and q (as a premise) are equivalent iff they make the same contribution to good implication.

As it turns out, the notions I've introduced so far allow us to define an intricate inferential role semantics. We need only enrich it with one additional operation, which I call adjunction. I shall now introduce the formal details. The semantics gives us a natural and intuitive way to think about the meaning of a sentence. Further, I produce fairly straightforward semantic clauses for connectives and these turn out to correspond to standard sequent calculus presentations of those connectives.[14]

2.1.2 The Formal Details

Definition 2.1.1 (Inferential Space $\mathbf{P}$, and Good Implications $\mathbb{I}$)**.** Let $\mathcal{L}$ be our language (of potential logical complexity) For my purposes here $\mathcal{L}$ is a propositional language, but there are natural extensions to first-order languages. An **inferential space** is the set of all ordered pairs of *multi-sets* of $\mathcal{L}$: $\mathbf{P} = \mathcal{P}(\mathcal{L})^2$. We call each "point" (of the form $\langle X, Y\rangle$, where $X, Y \subseteq \mathcal{L}$) an **implication**. Each inferential space $\mathbf{P}$ comes with a privileged subset of implications: the **good implications**: $\mathbb{I} \subseteq \mathbf{P}$.

Definition 2.1.2 (Adjunction)**.** There is a single associative and commutative operation on $\mathbf{P}$ called **adjunction**, '$\sqcup$'. If $A = \langle \Gamma, \Theta\rangle$ and $B = \langle \Delta, \Lambda\rangle$, then

$$A \sqcup B =_{df.} \langle \Gamma \cup \Delta, \Theta \cup \Lambda\rangle.$$

We also generalize '$\sqcup$' as an operation over subsets of $\mathbf{P}$.[15] If $X, Y \subseteq \mathbf{P}$, then:

$$X \sqcup Y = \{x \sqcup y | x \in X, y \in Y\}.$$

Next, I'll give a formal definition (aided by my recent definition of adjunction) of '$\curlyvee$' (pronounced "vee"). '$\curlyvee$' specifies the role that an ordered pair plays in good implication, i.e. which pairs in can be adjuncted such that the result is a good implication.

[14]The semantic picture below is heavily inspired by—in some sense isomorphic to a mere variation on—Girard's phase space space semantics. See (Girard, 2011).

[15]Recall that '$\cup$'—as in '$\Gamma \cup \Theta$'—is multi-set-union. Hence it may be that $\Gamma \cup \Gamma \neq \Gamma$). We treat the set of points as a *set* and hence subsets of $\mathbf{P}$ as sets of points—even if those points are structured by multi-sets.

Definition 2.1.3 (vee)**.** Suppose $X \subseteq \mathbf{P}$. Then:[16]

$$X^{\curlyvee} =_{df.} \{\langle \Delta, \Lambda \rangle \,|\, \forall \langle \Gamma, \Theta \rangle \in X\, (\langle \Gamma, \Theta \rangle \sqcup \langle \Delta, \Lambda \rangle \in \mathbb{I})\}.$$

With '$\curlyvee$' in hand, we can define when a set of implications is closed. Recall that earlier I said a second constraint on meaning is that contributions form equivalence classes of roles. This is because we wanted to fix extensionality of meaning at the level of contribution to inferential role. And it is easy to see that this defines a closure operation—since $(\cdot)^{\curlyvee}$ can be used to define a symmetric relation on subsets of $\mathbf{P}$, $(\cdot)^{\curlyvee\curlyvee}$ defines an equivalence relation.

Definition 2.1.4 (Closure)**.** A set of implications $X \subseteq \mathbf{P}$ is said to be **closed** iff $X^{\curlyvee\curlyvee} = X$.

Proposition 2.1.5. $(\cdot)^{\curlyvee\curlyvee}$ is a closure operation, i.e. $(\cdot)^{\curlyvee\curlyvee}$ is **extensive** ($X \subseteq X^{\curlyvee\curlyvee}$), **idempotent** ($X^{\curlyvee\curlyvee\curlyvee\curlyvee} = X^{\curlyvee\curlyvee}$) and **monotone** (if $X \subseteq Y$, then $X^{\curlyvee\curlyvee} \subseteq Y^{\curlyvee\curlyvee}$).

This closure operation allows us to introduce what I shall call "proper inferential roles" (PIRs). These are pairs of closed sets of implications. Proper inferential roles will serve as the interpretants of sentences. What is distinctive of *proper* inferential roles is that they satisfy both of the constraints that I rehearsed in the previous section.

Definition 2.1.6 (Proper Inferential Role)**.** A **proper inferential role (PIR)** is an ordered pair $\langle X, Y \rangle$ such that X and Y are each *closed*—in the sense defined above—subsets of $\mathbf{P}$ (i.e. $X^{\curlyvee\curlyvee} = X$ and $Y^{\curlyvee\curlyvee} = Y$).

When we think of $\langle X, Y \rangle$ as the semantic interpretant of a sentence A, we think of $\langle X, Y \rangle$ as consisting of a part, X, that specifies the role that A plays as a premise and a part, Y, that specifies the role A plays as a conclusion. The set of implications X are those implications into which A can be added as a premise to get a good implication. And Y is the set of implications into which A can be added as a conclusion to get a good implication. We say that X is the premissory role and Y the conclusory role of A.

[16] In the previous sub-section I defined '$\curlyvee$' over single ordered pairs. That definition, with the aid of adjunction would appear as follows (i.e. if X were a singleton—so suppose $X = \{\langle \Gamma, \Delta \rangle\}$):

$$X^{\curlyvee} = \{\langle \Gamma, \Delta \rangle\}^{\curlyvee} =_{df.} \{\langle \Delta, \Lambda \rangle | \langle \Gamma, \Delta \rangle \sqcup \langle \Delta, \Lambda \rangle \in \mathbb{I}\}.$$

Definition 2.1.7 (Convention)**.** As a convention if $[\![A]\!] = \langle X, Y\rangle$ is an inferential role, then we write $[\![A]\!]_P$ to refer to X and $[\![A]\!]_C$ to refer to Y, i.e. A's premissory and conclussory roles, respectively.

Philosophically, we should understand PIRs as specifying absolutely minimal, inferentialist constraints that semantic content must satisfy in accordance with Constraints One and Two above. In particular, in order to have specified the content of a sentence, (One:) one must specify both a premissory and a conclusory role for that content; and (Two:) those premissory and conclusory inferential roles must be closed, i.e., they must be sufficiently robust so that two roles cannot be merely notionally differentiated. For example suppose $X^{\curlyvee} = Y^{\curlyvee}$ (intuitively, X and Y play the same role in good implication), but $X \neq Y$. Then neither X nor Y is closed, since it is easy to show that $Y \subseteq X^{\curlyvee\curlyvee}$ and $X \subseteq Y^{\curlyvee\curlyvee}$.

Definition 2.1.8 (Models)**.** A **model** is a quadruple $\langle \mathcal{L}, \mathbf{P}, \mathbb{I}, [\![\cdot]\!]\rangle$ consisting of a language $\mathcal{L}$ and inferential space over that language $\mathbf{P}$, a privileged set of good implications $\mathbb{I}$, and an interpretation function $[\![\cdot]\!]$ (to be defined next) which interprets sentences in the language as inferential roles in the model.

Definition 2.1.9 (Interpretation Function)**.** An interpretation function $[\![\cdot]\!]$ maps sentences in $\mathcal{L}$ to proper inferential roles in models. If $A \in \mathcal{L}$ is atomic, then A is interpreted as follows:[17]

$$[\![A]\!] =_{df.} \langle \langle \{A\}, \emptyset\rangle^{\curlyvee}, \langle \emptyset, \{A\}\rangle^{\curlyvee}\rangle.$$

The semantic definitions of connectives follows:[18]

$$\begin{aligned}
[\![A\&B]\!] &=_{df.} \langle ((([\![A]\!]_P)^{\curlyvee} \sqcup ([\![B]\!]_P)^{\curlyvee})^{\curlyvee}, [\![A]\!]_C \cap [\![B]\!]_C\rangle,\\
[\![A \vee B]\!] &=_{df.} \langle [\![A]\!]_P \cap [\![B]\!]_P, ((([\![A]\!]_C)^{\curlyvee} \sqcup ([\![B]\!]_C)^{\curlyvee})^{\curlyvee}\rangle,\\
[\![A \to B]\!] &=_{df.} \langle [\![A]\!]_C \cap [\![B]\!]_P, ((([\![A]\!]_P)^{\curlyvee} \sqcup ([\![B]\!]_C)^{\curlyvee})^{\curlyvee}\rangle.,\\
[\![\neg A]\!] &=_{df.} \langle [\![A]\!]_C, [\![A]\!]_P\rangle.
\end{aligned}$$

Definition 2.1.10 (Semantic Entailment)**.** We say that A semantically entails B relative to a model $\mathcal{M}$ if the closure of the combination of A (as premise) and B (as conclusion)

[17]Note that, for arbitrary X, $X^{\curlyvee}$ is closed. Hence, we are guaranteed that $[\![A]\!]$ will be a PIR.

[18]I only introduce so-called "negative" connectives here, i.e. connectives with rules that are reversible.

consists of only good implications:

$$A \vDash_{\mathcal{M}} B \quad \text{iff}_{df.} \quad ((\llbracket A \rrbracket_P)^{\curlyvee} \sqcup (\llbracket B \rrbracket_C)^{\curlyvee})^{\curlyvee\curlyvee} \subseteq \mathbb{I}_{\mathcal{M}}.$$

We say that A semantically entails B if $A \vDash_{\mathcal{M}} B$ on all models $\mathcal{M}$.

NB: If A and B are multi-sets of sentences—suppose $A = \{A_1, \ldots, A_n\}$ and $B = \{B_1, \ldots, B_m\}$—then we read $A \vDash B$ as, for all models $\mathcal{M}$:

$$A_1, \ldots, A_n \vDash_{\mathcal{M}} B_1, \ldots, B_m \quad \text{iff}_{df.}$$
$$((\llbracket A_1 \rrbracket_P)^{\curlyvee} \sqcup \cdots \sqcup (\llbracket A_n \rrbracket_P)^{\curlyvee} \sqcup (\llbracket B_1 \rrbracket_C)^{\curlyvee} \sqcup \cdots \sqcup (\llbracket B_m \rrbracket_C)^{\curlyvee})^{\curlyvee\curlyvee} \subseteq \mathbb{I}_{\mathcal{M}}.$$

I will need to introduce more machinery, before I can prove anything interesting about entailment. But what I want to note here is the natural thought that what it means to say that A entails B (that B semantically follows from A) is nothing other than to say that the adjunction of A's premissory role and B's conclussory contains nothing but good implications. In the case where A and B are multi-sets, it is instead the adjunction of all premises (qua premises) and all conclusions (qua conclusions).

2.1.2.1 Generating Consequence Relations

In this section I explain how we may generate (substructural) consequence relations in line with the inferentialist view of meaning I sketched above.

Definition 2.1.11 (Base Consequence Relation (BCR))**.** A base consequence relation is a relation between finite multi-sets of atomic sentences, e.g. $\vdash_0 \subseteq \mathcal{P}(\mathcal{L}_0)^2$ (where $\mathcal{L}_0$ is the set of all atomic sentences of the language).

We start with what I call a "material" base consequence relation, i.e. a base consequence relation with no further constraints on it.

To be clear, I require none of the following constraints on a base consequence relation:

Definition 2.1.12 (Constraints on a base)**.** A base consequence relation may satisfy certain constraints.

Reflexive: A BCR is **reflexive** iff, for all $p \in \mathcal{L}_0$: $p \vdash_0 p$.

Axiom: If $\Gamma \vdash_0 \Theta$ then $\Gamma \vdash \Theta$.

$$\frac{\Gamma \vdash \Theta, A \qquad B, \Gamma \vdash \Theta}{A \to B, \Gamma \vdash \Theta}\ \text{L}\!\to \qquad \frac{A, \Gamma \vdash \Theta, B}{\Gamma \vdash A \to B, \Theta}\ \text{R}\!\to$$

$$\frac{\Gamma, A, B \vdash \Theta}{\Gamma, A\&B \vdash \Theta}\ \text{L\&} \qquad \frac{\Gamma \vdash \Theta, A \qquad \Gamma \vdash \Theta, B}{\Gamma \vdash \Theta, A\&B}\ \text{R\&}$$

$$\frac{A, \Gamma \vdash \Theta \qquad B, \Gamma \vdash \Theta}{A \vee B, \Gamma \vdash \Theta}\ \text{L}\vee \qquad \frac{\Gamma \vdash \Theta, A, B}{\Gamma \vdash \Theta, A \vee B}\ \text{R}\vee$$

$$\frac{\Gamma \vdash \Theta, A}{\neg A, \Gamma \vdash \Theta}\ \text{L}\neg \qquad \frac{A, \Gamma \vdash \Theta}{\Gamma \vdash \Theta, \neg A}\ \text{R}\neg$$

Figure 3: Sequent Calculus

Containment: A BCR satisfies **containment** iff, for all $p \in \mathcal{L}_0$ all $\langle \Delta, \Lambda \rangle \in \mathcal{P}(\mathcal{L}_0)^2$ we have: $\Delta, p \vdash_0 p, \Lambda$.

Monotonicity: A BCR is **monotonic** iff for all $\langle \Gamma, \Theta \rangle$ and $\langle \Delta, \Lambda \rangle$ in $\mathcal{P}(\mathcal{L}_0)^2$ we have: if $\Gamma \vdash_0 \Theta$ then $\Delta, \Gamma \vdash_0 \Theta, \Lambda$.

Contractive: A BCR is **contractive** iff $A, A, \Gamma \vdash_0 \Theta$ only if $A, \Gamma \vdash_0 \Theta$ and $\Gamma \vdash_0 \Theta, A, A$ only if $\Gamma \vdash_0 \Theta, A$ (for arbitrary Γ, Θ, A).

Transitive: A BCR is **transitive** iff $A, \Gamma \vdash_0 \Theta$ and $\Gamma \vdash_0 \Theta, A$ only if $\Gamma \vdash_0 \Theta$ (for arbitrary Γ, Θ, A.[19]

Next, I define a consequence relation, $\vdash$ as any sequent derivable from the sequent calculus in Figure 3, where the leaves are generated from the following single axiom (given a particular BCR $\vdash_0$).

2.1.2.2 Soundness and Completeness

Since we haven't imposed any restriction on models, however, it is easy to find counter-models for our base consequence relation. Thus, we need to appropriately limit models based

[19] We can define a similar version for mixed-cut. I do not in this document.

upon what $\vdash_0$ looks like.

Definition 2.1.13 (Base Consequence Relation)**.** A **base consequence relation** is a subset of $\mathbf{P}$ that consists of only atoms. B is a base consequence relation iff $B \subseteq \mathbf{P}$ and $B \cap \mathcal{P}(\mathcal{L}_0)^2 = B$.

We say that a model $\mathcal{M} = \langle \mathbf{P}, \mathbb{I}, [\![\cdot]\!] \rangle$ is **fit for** a base consequence relation B iff

$$\forall \langle \Delta, \Lambda \rangle \in B(\Delta \vDash_{\mathcal{M}} \Lambda).$$

We say that Γ semantically entails Θ **relative to** B iff $\Gamma \vDash_{\mathcal{M}} \Theta$ for all models $\mathcal{M}$ that are fit for B. We write this as $\Gamma \vDash_B \Theta$.

We are now in a position to show that the sequent calculus introduced in the previous section is sound and complete with respect to the inferential role semantics.

Theorem 2.1.14 (Soundness)**.** *The sequent calculus is sound:*

$$\Gamma \vdash_B \Theta \Rightarrow \Gamma \vDash_B \Theta.$$

Proof. The proof proceeds via induction on proof height. The base case is guaranteed by our restriction to models fit for B. To illustrate the proof, we do two interesting cases; the other cases are analogous. If the last step in our proof tree is R& (and so our sequent is $\Gamma \vdash \Theta, A\&B$), then we have $\Gamma \vDash_B \Theta, A$ and $\Gamma \vDash_B \Theta, B$, i.e.:

$$(([\![\&\Gamma]\!]_P)^\curlyvee \sqcup ([\![\bigvee(\Theta \cup \{A\})]\!]_C)^\curlyvee)^{\curlyvee\curlyvee} \subseteq \mathbb{I}$$

and

$$(([\![\&\Gamma]\!]_P)^\curlyvee \sqcup ([\![\bigvee(\Theta \cup \{B\})]\!]_C)^\curlyvee)^{\curlyvee\curlyvee} \subseteq \mathbb{I}.$$

So:

$$(([\![\&\Gamma]\!]_P)^\curlyvee \sqcup ([\![\bigvee(\Theta \cup \{A\})]\!]_C)^\curlyvee)^{\curlyvee\curlyvee} \cap (([\![\&\Gamma]\!]_P)^\curlyvee \sqcup ([\![\bigvee(\Theta \cup \{B\})]\!]_C)^\curlyvee)^{\curlyvee\curlyvee} \subseteq \mathbb{I}.$$

Since this is just the semantic definition of '&' qua conclusion, it is easy to see that we get the required semantic entailment:

$$(([\![\&\Gamma]\!]_P)^\curlyvee \sqcup ([\![\bigvee(\Theta \cup \{A\&B\})]\!]_C)^\curlyvee)^{\curlyvee\curlyvee} \subseteq \mathbb{I},$$

so:

$$\Gamma \vDash_B \Theta, A\&B.$$

Likewise, suppose the last step is R$\rightarrow$ (and so our sequent is $\Gamma \vdash \Theta, A \rightarrow B$). Then we have:

$$((\llbracket \&(\Gamma \cup \{A\}\rrbracket_P)^{\curlyvee} \sqcup (\llbracket \bigvee(\Theta \cup \{B\})\rrbracket_C)^{\curlyvee})^{\curlyvee\curlyvee} \subseteq \mathbb{I}.$$

Via some simple manipulations we reason:

$$((\llbracket \&\Gamma \rrbracket_P)^{\curlyvee} \sqcup (\llbracket \bigvee\Theta \rrbracket_C)^{\curlyvee} \sqcup (\llbracket A \rrbracket_P)^{\curlyvee} \sqcup (\llbracket B \rrbracket_C)^{\curlyvee})^{\curlyvee\curlyvee} \subseteq \mathbb{I}.$$

From which it is straightforward to show that

$$((\llbracket \&\Gamma \rrbracket_P)^{\curlyvee} \sqcup (\llbracket \bigvee\Theta \rrbracket_C)^{\curlyvee} \sqcup (\llbracket A \rightarrow B \rrbracket_C)^{\curlyvee})^{\curlyvee\curlyvee} \subseteq \mathbb{I}.$$

And thus:

$$((\llbracket \&\Gamma \rrbracket_P)^{\curlyvee} \sqcup (\llbracket \bigvee\Theta \cup \{A \rightarrow B\} \rrbracket_C)^{\curlyvee})^{\curlyvee\curlyvee} \subseteq \mathbb{I},$$

which just means

$$\Gamma \vDash_{\mathcal{M}} \Theta, A \rightarrow B.$$

■

Theorem 2.1.15 (Completeness)**.** *The sequent calculus is complete:*

$$\Gamma \vDash_B \Theta \Rightarrow \Gamma \vdash_B \Theta.$$

Proof. We show the contrapositive (i.e. that $\Gamma \nvdash_B \Theta \Rightarrow \Gamma \nvDash_B \Theta$) by constructing canonical models that can serve as counter-models for $\Gamma \nvdash_B \Theta$. To do this, we construct a model $\mathcal{M}$ which has the feature that $\Gamma \vDash_{\mathcal{M}} \Theta \Leftrightarrow \Gamma \vdash_B \Theta$.

It is easy to show (by induction on proof height) that we get this result if we define $\mathbb{I}_{\mathcal{M}}$ on our model as:

$$\mathbb{I}_{\mathcal{M}} =_{df.} \{\langle \Gamma, \Theta \rangle \in \mathbf{P} \mid \Gamma \vdash_B \Theta\}.$$

■

2.2 Representation Theorem & Logical Expressivism

In the previous section I produced a semantics for sentences which stand in (potentially radically) substructural relations of implication. That is, implications which may be non-monotonic, non-transitive, non-contractive, and non-reflexive. I also showed that the semantics is relatively tractable, insofar as relatively straightforward definitions of logical connectives allowed Gc3p (or Ketonen) style rules to be sound and complete for the semantics. This section compiles more evidence in favor of the claim that the semantics (and sequent calculus which is sound and complete with respect to it) is tractable and inherently interesting.

To do this, I prove a representation theorem for base consequence relations. That is, a method to get any theory (provided some modest constraints are met) via restrictions on base consequence relations. This result also gives us a way of precisifying a notion of some interest in the literature on inferentialism: namely logical expressivism.[20]

2.2.1 A Precisification of Logical Expressivism

I now seek to make the notion of "expression" more precise. Brandom understands expressivism in terms of what he calls an "LX relation", where a vocabulary B is "LX" of a vocabulary A if it is el̲aborated from and ex̲plicative of A. The first criterion (elaboration) has it that if one is able to successfully deploy vocabulary A then one already has the skills necessary to use B. That is, that B may be (algorithmetically) elaborated from A. The second criterion (explication) has it that B says something about (makes perspicuous in the object language) what one was doing by using A (minimally that B may encode the implications and incompatibilities of A). Logical vocabulary is said to be "universally LX" meaning that logical vocabulary stands in this relation to *all vocabularies*.

Let us make this relation more precise. First let $\mathcal{L}_0$ be an arbitrary vocabulary devoid of logical symbols (i.e. a set of atomic sentence letters). Let $\succ_0$ be a consequence relation over

[20] See Brandom (2008) for an original formulation. Peregrin (2014); Brandom (2018a,b) for some developments. See also Hlobil (2016) and Kaplan (2018). The latter of which is a predecessor to work carried out in the present section. Unpublished work by Ulf Hlobil, Ryan Simonelli, and Shuhei Shimamura has developed some of these themes.

$\mathcal{L}_0$ (i.e. $\mathrel{\succ\!\!-}_0 \subseteq \mathcal{P}(\mathcal{L}_0)^2$). Note that while I call $\mathrel{\succ\!\!-}_0$ and $\mathrel{\succ\!\!-}$ (below) *consequence* relations I do not yet impose *any* restrictions on them.[21] They should be treated, therefore, simply as two place relations between sets of sentences. As I discussed in the introduction, there are philosophically rich reasons for wanting a consequence relation that is e.g. non-monotonic or perhaps non-classical. In addition part of the motivation of expressivism is that where such features hold of consequence it is an *expression* of an underlying (material) relation of consequence.[22]

Next let $\mathbb{L}$ be our logic. Our logic consists of a finite set of logical symbols (e.g. $\{\&, \vee, \neg, \rightarrow\}$) and rules for expanding $\mathcal{L}_0$ to $\mathcal{L}$ (our language enriched with those logical symbols) and for expanding $\mathrel{\succ\!\!-}_0$ to $\mathrel{\succ\!\!-}$. Intuitively, we should think of $\mathbb{L}$ as a function from $\mathrel{\succ\!\!-}_0$ to $\mathrel{\succ\!\!-}$. That is, $\mathbb{L} : \mathrel{\succ\!\!-}_0 \mapsto \mathrel{\succ\!\!-}$. Then whether $\mathbb{L}$ is "LX" concerns the relationship between $\mathrel{\succ\!\!-}_0$ and $\mathrel{\succ\!\!-}$ (i.e. the behavior of $\mathcal{L}$ in relation to the behavior of $\mathcal{L}_0$).

That the logical vocabulary be elaborated fixes a tight relationship from $\mathrel{\succ\!\!-}_0$ to $\mathrel{\succ\!\!-}$. That is, to get from $\mathrel{\succ\!\!-}_0$ to $\mathrel{\succ\!\!-}$ should require no more than a specification of the logical vocabulary. That is, given $\mathrel{\succ\!\!-}_0$, $\mathrel{\succ\!\!-}$ should be uniquely determined: $\mathrel{\succ\!\!-}_0 \Rightarrow \mathrel{\succ\!\!-}$. In prose, the behavior of $\mathcal{L}$ should be determined by the behavior of $\mathcal{L}_0$ simply by specifying the logical symbols.

That the logical vocabulary be explicative fixes a tight relationship from $\mathrel{\succ\!\!-}$ to $\mathrel{\succ\!\!-}_0$. Since this requires that the logical vocabulary enable us to *say something* about the underlying pre-logical consequence relation, we should require that it actually do what it purports to do: $\mathrel{\succ\!\!-} \Rightarrow \mathrel{\succ\!\!-}_0$. In prose, the behavior of $\mathcal{L}$ should *genuinely* say or *express* something about the behavior of $\mathrel{\succ\!\!-}_0$. The behavior of $\mathcal{L}$ should therefore fix the behavior of $\mathcal{L}_0$. If $\mathcal{L}$ behaved differently then it would *express* something different about the behavior of $\mathcal{L}_0$. If such expression is to be genuine then the behavior of $\mathcal{L}_0$ (i.e. $\mathrel{\succ\!\!-}_0$) would need to be different.

Together these two criterion have it that $\mathrel{\succ\!\!-}_0 \Leftrightarrow \mathrel{\succ\!\!-}$. The behavior of $\mathcal{L}$ is elaborated out of, but also explicative of the behavior of $\mathcal{L}_0$.

[21] When speaking of consequence relations in a general sense, I will use "$\mathrel{\succ\!\!-}$" to formulate things. This is to make it clear that the formulation is meant more generally (i.e. it needn't be about the notion of consequence explored earlier in this chapter. I believe Lloyd Humberstone is responsible for this convention.

[22] I am being brief here on the *justification* for treating $\mathrel{\succ\!\!-}_0$ as I do. I primarily wish to stress here—in order to avoid confusion—that $\mathrel{\succ\!\!-}_0$ need not have *any constraints*.

While this criterion has some naive plausibility, it must still be made more precise. In particular, if our logical vocabulary is to be *conservative*, then $\mathrel{|\!\sim}_0 \subseteq \mathrel{|\!\sim}$ and so the criterion will hold trivially. We may circumvent this problem by quantifying over possible $\mathrel{|\!\sim}_0$. This might have already been anticipated since I mentioned that logical vocabulary is to have this relationship *universally* i.e. with respect to arbitrary vocabularies (and thus arbitrary $\mathrel{|\!\sim}_0$). This gives rise to the following definition:

Definition 2.2.1. Fix a logic $\mathbb{L}$, i.e. a function from $\mathrel{|\!\sim}_0$ to $\mathrel{|\!\sim}$. We say that $\mathbb{L}$ is **expressive** or that $\mathrel{|\!\sim}$ **expresses** a base consequence relation $\mathrel{|\!\sim}_0$ iff:

$$\begin{gathered}
(\forall \Gamma, \Theta \subseteq \mathcal{L})(\exists \Gamma_1, \Theta_1, \ldots, \Gamma_n, \Theta_n \subseteq \mathcal{L}_0) \\
(\forall \mathrel{|\!\sim}_0 \subseteq \mathcal{P}(\mathcal{L}_0)^2)(\forall \mathrel{|\!\sim} \subseteq \mathcal{P}(\mathcal{L})^2(\mathbb{L} : \mathrel{|\!\sim}_0 \mapsto \mathrel{|\!\sim})) \\
((\Gamma \mathrel{|\!\sim} \Theta) \Leftrightarrow (\Gamma_1 \mathrel{|\!\sim}_0 \Theta_1 \bigwedge \Gamma_2 \mathrel{|\!\sim}_0 \Theta_2 \bigwedge \cdots \bigwedge \Gamma_n \mathrel{|\!\sim}_0 \Theta_n)).
\end{gathered}$$

We also say $\Gamma \mathrel{|\!\sim} \Theta$ **expresses** $\Gamma_i \mathrel{|\!\sim}_0 \Theta_i$ $(1 \leq i \leq n)$ (its expressientia).

This definition says anytime $\Gamma \mathrel{|\!\sim} \Theta$ this is in virtue of some set of implications present in the language prior to $\mathbb{L}$. So the logical vocabulary is said to be elaborated if $\Gamma \mathrel{|\!\sim} \Theta$ occurs whenever those pre-logical implications obtain, and the logical vocabulary is said to be explicative if $\Gamma \mathrel{|\!\sim} \Theta$ occurs only if those pre-logical implications obtain.

The above should be taken as a precise specification of a minimal constraint on logics to count as "expressive". But one of the central features of expression is the idea that logical vocabulary should be able to make perspicuous in the object language *structural features* of inference. By structural features I have in mind such things as monotonicity, transitivity, contraction, reflexivity, classicality, etc., where each is understood to be capable of holding both globally (e.g. that $\mathrel{|\!\sim}$ is monotonic) as well as locally (e.g. that $\Gamma \mathrel{|\!\sim} \Theta$ is monotonic, though $\Delta \mathrel{|\!\sim} \Lambda$ may not be). Expressivism says that it is distinctive of logical vocabulary to be able to express such features. This requires that (i) $\mathbb{L}$ be capable of *preserving* structural features and that (ii) $\mathbb{L}$ be capable of *expressing* those very structural features it preserves.

Definition 2.2.2. Let $\mathfrak{Sf}$ be a structural feature. Officially, a structural feature is a property of elements of consequence relations. So we can think of $\mathfrak{Sf}$ as a family of partitions, one for each member of the family of consequence relations (i.e. a partition for each subset of

$\mathcal{P}(\mathcal{L}) \times \mathcal{P}(\mathcal{L})$).[23] Let $\mathfrak{Sf}(\Gamma \mathrel{\succ\!\!-} \Theta)$ be shorthand for $\Gamma \mathrel{\succ\!\!-} \Theta$ obeys (is an instance of) $\mathfrak{Sf}$. Next, let $\Gamma \mathrel{\succ\!\!-} \Theta$ be arbitrary with $\Gamma_i \mathrel{\succ\!\!-}_0 \Theta_i$ $(1 \leq i \leq n)$ its expressientia (in accordance with Definition 2.2.1). We say that a logic $\mathbb{L}$ **preserves** a structural feature $\mathfrak{Sf}$ iff:

$$\mathfrak{Sf}(\Gamma \mathrel{\succ\!\!-} \Theta) \Leftrightarrow \left(\mathfrak{Sf}(\Gamma_1 \mathrel{\succ\!\!-}_0 \Theta_1) \bigwedge \cdots \bigwedge \mathfrak{Sf}(\Gamma_n \mathrel{\succ\!\!-}_0 \Theta_n)\right).$$

A structural feature is *preserved* when an implication obeys that structural feature *iff* all of the implications it expresses also obey that structural feature. Thus, whether an implication obeys a structural feature should be seen as expressing something about the pre-logical implications that that implication expresses: it inherits those features from them and has those features in virtue of those implications alone. Next, I must explain what it means for a particular piece of logical vocabulary to express such structural features.

Definition 2.2.3. Let $\mathfrak{Sf}$ be a structural feature. Suppose some logical operation '*' may be used to mark a sequent in some way (with the constraint that $\Gamma^* \mathrel{\succ\!\!-} \Theta^*$ only if $\Gamma \mathrel{\succ\!\!-} \Theta$). Then we say that '*' (or $\mathbb{L}$) expresses $\mathfrak{Sf}$ iff there exists a '*' in $\mathbb{L}$ such that:

$$\Gamma^* \mathrel{\succ\!\!-} \Theta^* \Leftrightarrow \mathfrak{Sf}(\Gamma \mathrel{\succ\!\!-} \Theta).$$

Sf-Expression combines three ideas. (i) That a logic be capable of expressing an underlying base consequence relation, (ii) that it be capable of preserving structural features of that base consequence relation, and finally (iii) that it be able to mark in the object language those very same features that it preserves.

[23]Unofficially, a structural feature usually has something like an intension associated with it. It may be that the monotonic and transitive consequences of a consequence relation are the same. Nevertheless, the characterization of each is different, because the extensions may differ on a different consequence relation.

2.2.2 Representation Theorem

The logic I defined in the previous section is precisely such a logic. I repeat the details here for ease and prove some novel results surrounding the sequent calculus. Following this I show that the logic is expressive in the sense defined above by proving two representation theorems. Let us fix a base consequence relation (BCR) $\vdash_0$. Our logic include the symbols $\{\&, \vee, \neg, \rightarrow\}$ and expands $\mathcal{L}_0$ to $\mathcal{L}$ in the standard fashion. Then our logic $\mathbb{L}$ is given by the following sequent calculus, where proof trees are introduced by axioms (this is simply a reprinting of Figure 3 above):[24]

Axiom 1: If $\Gamma \vdash_0 \Theta$, then $\Gamma \vdash \Theta$ may form the base of a proof tree.

$$\frac{\Gamma \vdash \Theta, A \qquad B, \Gamma \vdash \Theta}{A \rightarrow B, \Gamma \vdash \Theta}\,\text{L}\!\rightarrow \qquad\qquad \frac{A, \Gamma \vdash \Theta, B}{\Gamma \vdash \Theta, A \rightarrow B}\,\text{R}\!\rightarrow$$

$$\frac{\Gamma, A, B \vdash \Theta}{\Gamma, A\&B \vdash \Theta}\,\text{L}\& \qquad\qquad \frac{\Gamma \vdash A, \Theta \qquad \Gamma \vdash B, \Theta}{\Gamma \vdash A\&B, \Theta}\,\text{R}\&$$

$$\frac{A, \Gamma \vdash \Theta \qquad B, \Gamma \vdash \Theta}{A \vee B, \Gamma \vdash \Theta}\,\text{L}\vee \qquad\qquad \frac{\Gamma \vdash A, B, \Theta}{\Gamma \vdash A \vee B, \Theta}\,\text{R}\vee$$

$$\frac{\Gamma \vdash A, \Theta}{\neg A, \Gamma \vdash \Theta}\,\text{L}\neg \qquad\qquad \frac{A, \Gamma \vdash \Theta}{\Gamma \vdash \neg A, \Theta}\,\text{R}\neg$$

Note that $\vdash_0$ and $\vdash$ here relate multisets. I treat things in this manner in order to avoid assuming *any* structural features absent permutation. I call $\mathbb{L}$ here NM-MS since its consequence relation is given by a Non-Monotonic Multi-Succedent sequent calculus.

Next I rehearse some important results for NM-MS .[25]

Theorem 2.2.4. *If $\Gamma \vdash \Theta$ may be arbitrarily weakened with atoms, then it may be arbitrarily weakened with logically complex sentences:*

$$\forall \Delta_0, \Lambda_0 \subseteq \mathcal{L}_0(\Delta_0, \Gamma \mid\!\sim \Theta, \Lambda_0) \Leftrightarrow \forall \Delta, \Lambda \subseteq \mathcal{L}(\Delta, \Gamma \mid\!\sim \Theta, \Lambda).$$

[24]The rules are the same as Ketonen uses. The rules with two top sequents are additive; the rules with a single top-sequent are multiplicative. These are sometimes called "mixed" or "assorted" rules/connectives (see e.g. Dicher, 2016). It is similar to the system called G3cp discussed in (Negri et al., 2008, ch. 3) with a more standard treatment of negation and material axioms. As is well known, these rules are equivalent to the multiplicative and additive rules of linear logic given monotonicity and contraction (Girard, 2011).

[25]Note that many of these results have full proofs worked out in (Girard, 2011; Negri et al., 2008). Since my system is slightly different than the systems featured there, a more thorough treatment would require some minor modification.

Proof. ($\Leftarrow$) is immediate. ($\Rightarrow$) is proven by induction on the complexity of $\Delta \cup \Lambda$ where complexity is understood in terms of the complexity of the most complex sentences in $\Delta \cup \Lambda$. ■

A similar result is in the offing, namely that the sequent calculus preserves contraction.

Theorem 2.2.5. *If* $\Gamma \mathrel{|\!\!\sim} \Theta$ *allows contraction of atomic sentences, then it allows contraction of logically complex sentences.*

Proof. One direction is trivial, the other direction is provided by induction on the complexity of the contracted sentence. ■

Since it is well known that the rules featured above are equivalent to both the additive and multiplicative rules of linear logic given contraction and monotonicity, we can actually locate the condition needed for our logic to be supra-classical.

Definition 2.2.6. We say that $\vdash_0$ obeys Containment (CO) if

$$\forall \Delta, \Lambda \subseteq \mathcal{L}_0 (\Delta, p \vdash_0 p, \Lambda)$$

(i.e. if we have $\forall q \in \mathcal{L}_0 (q \vdash_0 q)$ and all such sequents may be arbitrarily weakened; the fragment carved out by this stipulation will also obviously obey contraction). In short: let us define $\vdash_0^{CO}$ such that $\vdash_0^{CO}$ obeys reflexivity $\forall q \in \mathcal{L}_0 (q \vdash_0 q)$, weakening and contraction. And further stipulate that no proper subset of $\vdash_0^{CO}$ obeys all of these conditions. A base consequence relation $\vdash_0$ is said to obey CO iff it includes $\vdash_0^{CO}$, i.e. $\vdash_0^{CO} \subseteq \vdash_0$.

Theorem 2.2.7. *If* $\vdash_0$ *obeys CO, then* $\vdash$ *is supra-classical.*

Proof. Result is well known, but can be easily established by showing an equivalence with Gentzen's LK in the fragment of $\vdash$ generated by $\vdash_0^{CO}$. ■

Finally, the next theorem is of particular import to the sections following this one.

Theorem 2.2.8. *All rules of the sequent calculus are reversible. That is, if* $\Gamma \vdash \Theta$ *would be the result of the application of a rule to* $\Gamma^* \vdash \Theta^*$ *(and possibly* $\Gamma^{**} \vdash \Theta^{**}$*) then*

$$\Gamma \vdash \Theta \Leftrightarrow \Gamma^* \vdash \Theta^* (and \quad \Gamma^{**} \vdash \Theta^{**}).$$

Proof. Proof is straightforward by induction on proof height. ∎

From this my first gloss on logical expression follows immediately. In the next section I prove that the more precise sense (in Definition 2.2.1) also holds.

Corollary 2.2.9. *The sequent calculus is conservative. That is*

$$\Gamma \vdash_0 \Theta \Leftrightarrow \Gamma \vdash \Theta.$$

Next I show how consequence relations may be represented in NM-MS. First two central results concerning conjunctive and disjunctive normal forms.[26]

Proposition 2.2.10. Let $CNF(A)$ be the conjunctive normal form representation of A. It follows that

$$\Gamma \vdash \Theta, A \Leftrightarrow \Gamma \vdash \Theta, CNF(A).$$

Proof. Proof proceeds constructively. From theorem 2.2.8, we may deconstruct A until we have a number of sequents of the form: $\Gamma \vdash \Theta, \Lambda_1$; $\Gamma \vdash \Theta, \Lambda_2; \ldots$ $\Gamma \vdash \Theta, \Lambda_n$ where $\Lambda_i (1 \leq i \leq n)$ contains only literals. We next construct $CNF(A)$ via repeated application of R$\vee$ and R&:

$$\Gamma \vdash \Theta, (\bigvee \Lambda_1) \& (\bigvee \Lambda_2) \& \cdots \& (\bigvee \Lambda_n),$$

i.e. $\Gamma \vdash \Theta, CNF(A)$. ∎

Proposition 2.2.11. Let $DNF(A)$ be the disjunctive normal form representation of A. It follows that

$$A, \Gamma \vdash \Theta \Leftrightarrow DNF(A), \Gamma \vdash \Theta.$$

Proof. Proof is identical to the previous proposition except the sets are on the left and we construct $DNF(A)$ via L& and L$\vee$. ∎

Theorem 2.2.12 (Representation Theorem 1)**.** *Let CR be a consequence relation, i.e. $CR \subseteq \mathcal{P}(\mathcal{L})^2$. Then we may specify what must be included in $\vdash_0$ such that $CR \subseteq \vdash$.*

[26]Note that the results in Propositions 2.2.10 and 2.2.11 follow closely the distribution properties Girard demonstrates for different connectives in linear logic (Girard, 1987, 2011).

Proof. Proof proceeds constructively. For each $\Gamma \vdash \Theta$ in CR let us find an equivalent $CNF(A) \vdash CNF(B)$. This has the form:

$$(\&\Gamma_1) \vee \cdots \vee (\&\Gamma_a) \vdash (\bigvee\Theta_1)\& \cdots \&(\bigvee\Theta_b).$$

This holds just in case (for $1 \leq i \leq a$ and $1 \leq j \leq b$) $\Gamma_i \vdash_0 \Theta_j$. Thus we stipulate of the base that $\Gamma_i \vdash_0 \Theta_j$ for $1 \leq i \leq a$ and $1 \leq j \leq b$. If we do this for each implication in CR then we are guaranteed that $CR \subseteq\vdash$. ■

Theorem 2.2.13 (Representation Theorem 2)**.** *Let CR be a consequence relation. If CR is closed under some modest syntactic constraints,*[27] *then we may specify $\vdash_0$ such that $CR =\vdash$.*

Proof. Proof is identical to the first Representation Theorem except that the syntactic constraints on CR have it that $\mid\!\sim = CR$. ■

These results give us a way of saying exactly how to reconstruct arbitrary consequence relations using my machinery and given some modest constraints how to reconstruct them *exactly*. It is this ability to reconstruct consequence relations *exactly* that will prove most important. For what it shows is that we are able to find exactly which pre-logical implications an arbitrary implication involving logical vocabulary *expresses*. That is, what I have shown is a method for finding exactly which implications in $\vdash_0$ are expressed by each implication in $\vdash$. We are thus in a position to prove the following straight away.

Theorem 2.2.14 (Expressivity)**.** *NM-MS is expressive. That is, we have*

$$\Gamma \mid\!\sim \Theta \Leftrightarrow (\Gamma_1 \mid\!\sim_0 \Theta_1 \bigwedge \cdots \bigwedge \Gamma_n \mid\!\sim_0 \Theta_n).$$

for some $\Gamma_1, \Theta_1, \ldots, \Gamma_n, \Theta_n$ and arbitrary $\mid\!\sim_0$.

[27] In a more formal account I treat representation as of *theories* (see the appendix). Here I characterize it in terms of consequence relations, where we are able to precisely represent a consequence relation just in case it is closed under the rules of NM-MS. In the case where we wish to treat theories instead, then a theory T must meet the following constraints: &-composition and -decomposition ($A, B \in T$ iff $A\&B \in T$), Distribution (of $\vee$ over &) $(A \vee (B\&C)) \in T$ iff $(A \vee B)\&(A \vee C) \in T$, Conditional Equivalence ($A \to B = \sigma$ is a sub-formula of $\tau \in T$ iff $\neg A \vee B = \sigma'$ is a subformula of $\tau \in T$), both De-Morgan's Equivalences (likewise defined over sub-formulae) and involution (also defined over subformulae).

Proof. Suppose $\Gamma \vdash \Theta$ and let it be equivalent to $DNF(A) \vdash CNF(B)$ for some A and B. This has the form:

$$(\&\Gamma_1) \vee \cdots \vee (\&\Gamma_a) \vdash (\bigvee\Theta_1)\& \cdots \&(\bigvee\Theta_b).$$

This holds just in case (for $1 \leq i \leq a$ and $1 \leq j \leq b$) $\Gamma_i \vdash_0 \Theta_j$. Next, let us enumerate $\langle i, j\rangle$ as $1, \ldots, n$. Then we have that:

$$\Gamma \vdash \Theta \Leftrightarrow (\Gamma_1 \vdash_0 \Theta_1 \bigwedge \cdots \bigwedge \Gamma_n \vdash_0 \Theta_n).$$

■

2.2.3 Recovering Structure: Sf-Expression

I have so far shown how NM-MS is expressive in the sense made precise in Definition 2.2.1. Now I show how NM-MS may express particular structural features. First I introduce a schema for introducing a piece of logical vocabulary '$\boxed{\mathfrak{S}}$'.

First, let us enrich our sequent calculus by introducing a second turnstile $\vdash^{\mathfrak{S}}$. Now let $\vdash_0^{\mathfrak{S}}$ pick out some subset of $\vdash_0$. Later I will discuss principles for determining which subset, but for now I leave the details vague. We may introduce the following rules to our sequent calculus:[28]

Axiom 2: If $\Gamma \vdash_0^{\mathfrak{S}} \Theta$ then $\Gamma \vdash^{\mathfrak{S}} \Theta$.

$$\frac{A, \Gamma \vdash^{\mathfrak{S}} \Theta}{\boxed{\mathfrak{S}}A, \Gamma \vdash^{[\mathfrak{S}]} \Theta}\ \mathrm{L}\boxed{\mathfrak{S}} \qquad\qquad \frac{\Gamma \vdash^{\mathfrak{S}} \Theta, A}{\Gamma \vdash^{[\mathfrak{S}]} \Theta, \boxed{\mathfrak{S}}A}\ \mathrm{R}\boxed{\mathfrak{S}}$$

Lemma 2.2.15. *L$\boxed{\mathfrak{S}}$ and R$\boxed{\mathfrak{S}}$ are reversible rules.*

We thus have the following result.

Theorem 2.2.16. *Let $\mathfrak{Sf}$ be a structural rule. Suppose that $\mathfrak{Sf}$ is* preserved *(in the sense of Definition 2.2.2) and suppose further that $\vdash^{\mathfrak{S}}$ marks that structural feature exactly. We*

[28]Note that the rest of our sequent calculus is altered such that our other rules preserve $\vdash^{\mathfrak{S}}$. E.g. R& requires that both top sequents have either $\vdash$ or $\vdash^{\mathfrak{S}}$ (I do not allow mixed cases).

thus have: $\mathfrak{S}\mathfrak{f}(\Gamma \vdash \Theta)$ *iff* $\Gamma \vdash^{\mathfrak{S}} \Theta$. *It follows that* $\boxed{\mathfrak{S}}$ expresses *(in the sense of Definition 2.2.3)* $\mathfrak{S}\mathfrak{f}$. *Thus:*

$$\boxed{\mathfrak{S}}A, \Gamma \vdash \Theta \Leftrightarrow \mathfrak{S}\mathfrak{f}(A, \Gamma \vdash \Theta)$$

$$\Gamma \vdash \Theta, \boxed{\mathfrak{S}}A \Leftrightarrow \mathfrak{S}\mathfrak{f}(A, \Gamma \vdash \Theta, A)$$

Proof. I prove only the latter biconditional since the proof of the former is identical. By supposition $\mathfrak{S}\mathfrak{f}(\Gamma \vdash \Theta, A)$ iff $\Gamma \vdash^{\mathfrak{S}} \Theta, A$. Since it follows that our R$\boxed{\mathfrak{S}}$ rule is reversible, we have that $\Gamma \vdash^{\mathfrak{S}} \Theta, A$ iff $\Gamma \vdash \Theta, \boxed{\mathfrak{S}}A$. Thus

$$\Gamma \vdash \Theta, \boxed{\mathfrak{S}}A \Leftrightarrow \mathfrak{S}\mathfrak{f}(\Gamma \vdash \Theta, A).$$

■

The result of the above proof is a general method for introducing logical vocabulary that is *expressive* of structural features. If the rules for the logical vocabulary's introduction are reversible and the structural feature in question is *preserved* by $\mathbb{L}$, then the logical vocabulary will *express* that structural feature. I next rehearse two specific cases of this: an operator that marks monotonicity and an operator that marks classical validity.

2.2.3.1 Expressing Structural Features

The rules for monotonicity have the following form:

Axiom 2: If $\forall \Delta, \Lambda \subseteq \mathcal{L}_0(\Delta, \Gamma \vdash_0 \Theta, \Lambda)$ then $\Gamma \vdash^M \Theta$.

$$\frac{A, \Gamma \vdash^M \Theta}{\boxed{M}A, \Gamma \vdash^{[M]} \Theta}\ \text{L}\boxed{M} \qquad\qquad \frac{\Gamma \vdash^M \Theta, A}{\Gamma \vdash^{[M]} \Theta, \boxed{M}A}\ \text{R}\boxed{M}$$

I have already shown in Theorem 2.2.4 that weakening is preserved by the rules of NM-MS. It therefore follows that:

Corollary 2.2.17. $\boxed{M}$ *expresses weakening/monotonicity. That is,*

$$\boxed{M}A, \Gamma \vdash \Theta \Leftrightarrow \forall \Delta, \Lambda(\Delta, A, \Gamma \vdash \Theta, \Lambda)$$

$$\Gamma \vdash \Theta, \boxed{M}A \Leftrightarrow \forall \Delta, \Lambda(\Delta, \Gamma \vdash \Theta, A, \Lambda)$$

This means that we may expand NM-MS (our $\mathbb{L}$) in order to mark *in the object language* which implications are persistent under arbitrary weakenings. Next, I show a similar result for contraction. That is, I show a way of marking sequents that are where contraction holds. We introduce the following axiom and rules as before:

Axiom 2: If $\Gamma \vdash_0 \Delta$ and for arbitrary Δ, Λ we have $\Delta \vdash_0 \Lambda$ if this would be the result of some number of applications of contraction to $\Gamma \vdash_0 \Delta$, then $\Gamma \vdash^C \Theta$.

$$\frac{A, \Gamma \vdash^C \Theta}{\boxed{C}A, \Gamma \vdash^{[C]} \Theta}\ \mathrm{L}\boxed{C} \qquad\qquad \frac{\Gamma \vdash^C \Theta, A}{\Gamma \vdash^{[C]} \Theta, \boxed{C}A}\ \mathrm{R}\boxed{C}$$

I have already shown in Theorem 2.2.5 that contraction is preserved by the rules of NM-MS. It therefore follows that:

Corollary 2.2.18. $\boxed{C}$ *expresses contraction. That is,*

Next, I demonstrate the same for "classicality", i.e. develop an operator that marks implications that are valid classically.

Axiom 2: If $\Gamma, p \vdash_0 p, \Theta$ then $\Gamma, p \vdash^K p, \Theta$ (where Γ, Θ may be possibly empty).

$$\frac{A, \Gamma \vdash^K \Theta}{\boxed{K}A, \Gamma \vdash^{[K]} \Theta}\ \mathrm{L}\boxed{K} \qquad\qquad \frac{\Gamma \vdash^K \Theta, A}{\Gamma \vdash^{[K]} \Theta, \boxed{K}A}\ \mathrm{R}\boxed{K}$$

Again, I have already shown in Theorem 2.2.7 that classicality is a feature NM-MS preserves. Thus any sequent which is derived from atomic sequents which are part of the CO (cf. Definition 2.2.6) fragment of $\vdash_0$ (regardless of whether $\vdash_0$ actually obeys CO) will be classically valid.

Corollary 2.2.19. *Let* $\vdash_{LK}$ *be the consequence relation instantiated by Gentzen's LK minus the rules for quantifiers (and with* $\wedge$ *substituted with &, etc.). Then* $\boxed{K}$ *expresses classical validity, that is:*

$$\boxed{K}A, \Gamma \vdash \Theta \Leftrightarrow A, \Gamma \vdash_{LK} \Theta$$

$$\Gamma \vdash \Theta, \boxed{K}A \Leftrightarrow \Gamma \vdash_{LK} \Theta, A$$

There are of course many further possibilities for such '$\boxed{\mathfrak{S}}$' operators. We may also introduce vocabulary for expressing inference that obey, transitivity + weakening, more restricted weakening principles, and perhaps more.[29]

I will introduce one more such notion. I introduce it because it will have an interesting philosophical usage in the fourth chapter of the dissertation (though I only make short reference to it there). That is a way of marking regions of the consequence relation which are:

- Supra-classical (contain CO)
- Monotonic
- Transitive
- Contractive

I claim that this combination of features captures intuitively the notion of a sentence's "literal meaning". What follows form a sentence, strictly speaking (or taken literally),[30] should follow from that sentence regardless of the context under which it is being considered (or what one uses the sentence to mean). Included within this are all of the classical consequences as well as consequences which are persistent under arbitrary weakenings. It also seems to me that sentences, taken literally, should contract (since we're isolating a fragment of the sentence meant to express something like a "minimal proposition").[31]

We introduce the following rules for "literally".

Axiom 2: It is simpler to define the anti-extension. Suppose $A, \Gamma \vdash_0 \Theta$. Then $A, \Gamma \not\vdash^L \Theta$ if any of the following hold:

- Exists $p \in \{A\} \cup \Gamma \cup \Theta$ and $\langle \Lambda, \Delta \rangle$ such that $\Lambda, p \not\vdash_0 p, \Delta$.
- $\Gamma \vdash_0 \Theta, A$ but $\Gamma \not\vdash_0 \Theta$
- $\Delta \not\vdash_0 \Theta$ where this would be the result of contracting the original sequent
- There exists $\langle \Lambda, \Delta \rangle$ such that $A, \Lambda, \Gamma \not\vdash_0 \Theta, \Delta$.

[29] Makinson for example considers a consequence relation which is supra-classical, monotonic, and obeys transitivity (Makinson, 1994, 2005). We could introduce an operator to express exactly this consequence relation along with some other consequence relations discussed therein.

[30] The locution "what follows strictly speaking" shouldn't be read to mean "strict implication". That is *not* what I mean to capture.

[31] The significance of contraction here might not be fully appreciated until the next chapter.

Similarly for $\Gamma \vdash_0 \Theta, A$.

$$\frac{A, \Gamma \vdash^{L} \Theta}{\boxed{L}A, \Gamma \vdash^{[L]} \Theta}\ \mathrm{L}\boxed{L} \qquad\qquad \frac{\Gamma \vdash^{L} \Theta, A}{\Gamma \vdash^{[L]} \Theta, \boxed{L}A}\ \mathrm{R}\boxed{L}$$

NM-MS already preserves contraction, weakening, and classicality. So we must show that transitivity is obeyed in that fragment as well. We show the result we want directly.

Theorem 2.2.20. *The following rule would be eliminable:*[32]

$$\frac{A, \Gamma \vdash^{L} \Theta \qquad \Gamma \vdash^{L} \Theta, A}{\Gamma \vdash^{L} \Theta}$$

Proof. It is sufficient to show we may push cut up the proof tree via induction on complexity of A. The base case is immediate via Axiom 2 and the reversibility of the rules. Here is how the inductive step works for conjunction: suppose A is of the form $A\&B$. Then the top sequents are:

$$A\&B, \Gamma \vdash^{L} \Theta,$$
$$\Gamma \vdash^{L} \Theta, A\&B.$$

Since the rules are reversible we have:

$$A, B, \Gamma \vdash^{L} \Theta,$$
$$\Gamma \vdash^{L} \Theta, A,$$
$$\Gamma \vdash^{L} \Theta, B.$$

Since monotonicity is preserved we are guaranteed: $B, \Gamma \vdash^{L} \Theta, A$. Via the inductive hypothesis we have:

$$B, \Gamma \vdash^{L} \Theta.$$

Via another invocation of the same we have our result:

$$\Gamma \vdash^{L} \Theta.$$

■

[32] Because this fragment is both monotonic and contractive I use a shared-context version of cut. We could just as well have used the rule:

$$\frac{A, \Gamma \vdash^{L} \Theta \qquad \Lambda \vdash^{L} \Delta, A}{\Gamma, \Lambda \vdash^{L} \Theta, \Delta}$$

Thus we have a way of marking the confluence of all of these structural features *in the object language*. I claim that this *expresses* "minimal propositions" (or the literal meaning) of a sentence. This does not mean that I endorse these ideas. But I do think that for sentences that seem to have them, that we can mark this behavior.[33]

2.2.3.2 Some Defective Cases

So far I have introduced a more precise criterion for understanding logical expressivism and in particular for understanding how *structural* features of inference might be expressed. I then introduced a system that was not only *expressive* in this sense, but also successfully preserved and expressed several important structural features. In order to appreciate exactly what I am up to, however, it will be useful to look at some cases where each of these criteria fail.

Example 2.2.21. The multiplicative rules of linear logic are *not expressive*. I show that this is the case for the multiplicative conjunction $\otimes$:

$$\frac{\Gamma, A, B \vdash \Theta}{\Gamma, A \otimes B \vdash \Theta}\,\text{L}\otimes \qquad\qquad \frac{\Gamma \vdash \Theta, A \qquad \Delta \vdash \Lambda, B}{\Gamma, \Delta \vdash \Theta, \Lambda, A \otimes B}\,\text{R}\otimes$$

It is sufficient to show a case where the logic does not express particular implications in $\vdash_0$. Notice that there are potentially two ways to derive $p \otimes q \vdash p \otimes q$ where $p, q \in \mathcal{L}_0$:

$$\frac{\dfrac{p \vdash p \qquad q \vdash q}{p, q \vdash p \otimes q}\,\text{R}\otimes}{p \otimes q \vdash p \otimes q}\,\text{L}\otimes \qquad\qquad \frac{\dfrac{p, q \vdash q}{p \otimes q \vdash q}\,\text{L}\otimes \qquad \vdash p}{p \otimes q \vdash p \otimes q}\,\text{R}\otimes$$

Since the atomic sequents used to start each proof tree are different (in fact they are entirely different), it's possible that $\vdash_0$ includes one and $\vdash_0'$ includes the other and thus the presence of $p \otimes q \vdash p \otimes q$ does not guarantee the presence of either. In this sense, logics which include '$\otimes$' are not expressive in the relevant sense.

It is also possible to find counter-examples to Sf-Preservation and Sf-Expression. Even using the rules of NM-MS such counter-examples will arise:

[33] Hopefully these remarks aren't too opaque. Proponents who endorse the idea of a literal meaning often take it that the assertion of "p" and the assertion of "p literally" (where literally is appropriately inserted) are identical. I certainly don't claim that $\boxed{L}p$ and p are the same.

Example 2.2.22. Suppose we want to introduce an operator '$\boxed{R}$' to mark instances of reflexivity, i.e. $\phi \vdash \phi$. Then the rules for introducing such an operator should probably have the form:

Axiom 2: If $p \vdash_0 p$ then $p \vdash^R p$.

$$\frac{A, \Gamma \vdash^R \Theta}{\boxed{R}A, \Gamma \vdash^{[R]} \Theta}\ \mathrm{L}\boxed{R} \qquad\qquad \frac{\Gamma \vdash^R \Theta, A}{\Gamma \vdash^{[R]} \Theta, \boxed{R}A}\ \mathrm{R}\boxed{R}$$

Unfortunately, it is easy to show that NM-MS fails to preserve reflexivity and thus fails to express it. For example $A\&B \vdash A\&B$ is clearly an instance of reflexivity and thus we should want $A\&B \vdash \boxed{R}(A\&B)$. But clearly $A\&B \vdash A\&B$ must be derived from $A, B \vdash A$ and $A, B \vdash B$, neither of which are instances of reflexivity.[34]

There will therefore be logics which in general fail to be expressive and even among those that are expressive there will be structural features that fail to be preserved and thus expressed. Deciding how expressive one wants one's logic to be and which structural features ought to be preserved are therefore *not* independent questions.

2.3 Why go substructural?

So far I have motivated a formal semantics based upon the idea that (i) the content of a sentence should be understood in terms of the contribution that that sentence makes to good implication and (ii) that such implications may be radically substructural (i.e. may disobey monotonicity, transitivity, contraction, and/or reflexivity in potentially unpredictable ways). I have also shown that despite the radically substructural nature of the implication relation that underlies this notion of content, that a perfectly tractable semantics emerges on which much work can be done, including a proof of a representation theorem. In the next chapter,

[34]Though they are both found in the region of the consequence relation which allows reflexivity *together with weakening*, hence why we are able to have an operator to mark classicality.

It is also worth remarking that the above might also fail for independent reasons. For example, if we are able to derive $A\&B \vdash A\&B$, then we could also derive $A\&B \vdash B\&A$, but is the latter here an instance of the structural feature of reflexivity? It is not obvious that we should think so. In general, even when a sequent calculus preserves reflexivity, it needn't generate *only* reflexive sequents from the reflexive fragment of its axioms.

I will try to get clearer on what makes this approach to semantics particularly compelling and why it has been so far under-explored (i.e. I intend to diagnose what has motivated an avoidance of this approach). Before I can do that, however, I want to put forward some brief arguments for why we should even go substructural (as opposed to accounting for this phenomena in some other manner).

The argument in this section are therefore meant to motivate and not decide the matter. In short: I've shown that you *can* construct such a semantics, now I want to show why you might *want to* (and next: why motivations against such a project are misguided). To do this, I'll examine one way of trying to account for substructural consequence that fails to bear fruit. My arguments here are meant to be philosophical, not technical. In a fuller treatment, I believe that similar arguments could be made for other structural features, but I only pursue one such argument here. In particular, I argue that attempts to understand non-monotonicity in terms of defeat or defeasible reasoning flounder. This is because the central notion required for such accounts—that of a "defeater"—cannot be made sense of. Instead, I hope to illustrate that when Γ implies A but some superset of Γ doesn't:

$$\Gamma \vdash A$$
$$\Gamma, \Delta \nvdash A.$$

It need not be because of some defeater in the latter case, but rather because of the way in which the various considerations in $\Gamma \cup \Delta$ hang together. To understand this phenomenon in terms of how these considerations hang together is precisely to understand it as a structural phenomenon, i.e. in terms of non-monotonicity rather than defeat.

First, I'll briefly rehearse how defeat has been understood in the epistemology and defeasible reasoning literatures (§2.3.1). Following this, I try to isolate a tractable notion of defeat with ever weaker conceptions of it (§2.3.2). But none of these are shown to work.

2.3.1 Understanding Defeaters: An Exercise in Defeat

The concept of defeat has figured prominently within epistemology since Gettier (1963). The counter-examples that Gettier employs against the justified true belief (JTB) account of knowledge often rely on deviant causal chains (or similar devices) and the thought was that a "no defeater" condition on knowledge might save the JTB account from Gettier style counter-examples.[35] If a defeater is a consideration whose absence begets knowledge, then a "no defeater" condition seems like a plausible way to avoid Gettier style counter-examples. The idea was apparently compelling enough that it generated its own subfield: that of *defeasible reasoning*.[36] Defeasible reasoning is reasoning which is subject to defeaters. That is, if A is a reason for B, then we can characterize this reason as *defeasible* if the presence of some further consideration C makes it such that A is no longer a reason for B.

This chapter seeks to raise challenges for the concept of defeat. While recently the place of the concept of defeat within epistemology has seen some challenges concerning its

[35]The actual history is at least one more degree more complicated. The concept finds its origin in the concepts of *defeasance* and *defeasibility*, which originated in English common law under the Tudors (Geisst, 2013; Plucknett, 1956)). A defeasance in this technical sense is a separate instrument that nullifies or cancels an interest in a property.

This concept apparently laid dormant in property law until, sensing a lacuna in the legal conceptual space, Hart (1948) appropriated the notion of defeasibility for contracts generally, and therefore to ascriptions of rights and responsibilities. In short: while it is common to speak of certain conditions needed for e.g. a contract, there was no term for things which, should they be present, would defeat or undo a contract.

It did not take long for the concept to find its way into other areas of philosophy. For example, Chisholm (1964) appropriates the idea of defeat, defeasibility, etc. to make sense of Ross' notion of a *prima facie* duty. While it is common to cash out *prima facie* and *pro tanto* duties in terms of defeat and defeasibility this language is totally lacking from Ross, whose writing predates Hart's appropriation. Chisholm also appears to be the first to use this language in epistemology in his (1966). While Chisholm's usage there and the development of research into defeasible reasoning was and is much more broad, it seems that it entered into epistemology at least in part in response to Gettier's challenge to JTB accounts of knowledge.

According to this dialectic, Gettier showed that JTB is insufficient for knowledge. Many responded by positing that a fourth condition was needed on knowledge: a no-Gettier-condition. Many of these fourth conditions seemed to amount to no defeater views (for example: Sosa (1964); Lehrer (1965); Pollock (1967); Sosa (1969)). Just a few years later—post Chisholm's appropriation into epistemology—similar accounts (from some of the same authors) offer a "no defeater" response to Gettier and explicitly use the language of defeat (e.g. Lehrer and Paxson Jr (1969); Sosa (1970); Lehrer (1970); Klein (1971); Swain (1972b,a)). This dialectic, and in particular the role of defeat in the literature surrounding Gettier is suggested explicitly by Swain (1974); Levy (1977). Lehrer (2019) ties the emergence of defeat in epistemology to Gettier explicitly.

This description, however, is perhaps not the most accurate narrative surrounding its emergence in epistemology. Rather, it seems that the emergence of the idea of defeat in epistemology as an answer to Gettier and Chisholm's appropriation more or less coincided. It is hard to think that such a coincidence was independent given the figures involved, but it would also be slightly distortive to suggest that Chisholm appropriated the notion *in order to* respond to Gettier; nevertheless, it wouldn't have emerged as it has (epitomized by e.g. Lehrer and Paxson Jr (1969)) had it not been for Gettier. Obviously, the concept was of independent interest and so was quickly theorized on its own and put to work in e.g. AI, rational choice, etc.. (see e.g. Pollock (1970, 1974)).

[36]See Koons (2017). As with the emergence of the concept of defeat, the emergence of the field of defeasible reasoning was also likely the result of convergence and coincidence. Around the same time that defeat emerged as a concept in epistemology, researchers in artificial intelligence needed a way to describe the kind of largely heuristical reasoning needed for reasoning outside of mathematics and logic (for example; McCarthy and Hayes (1969); see also Lehrer's retrospective mentioned above (Lehrer, 2019)).

importance,[37] my intent is to go one step further and to challenge the internal coherence of the concept. That is, I do not think that *defeat* is a unified phenomenon and rather acts like a grab bag for a number of different things. In particular I intend to attack what I will call a *defeatist* conception of reasoning. The *defeatist* is committed to the following:

Definition 2.3.1 (Defeatism)**.** It sometimes happens that some set of considerations Γ provide reason for A. Nevertheless, it may happen that Δ (where $\Delta \supsetneq \Gamma$) does *not* provide reason for A. In such a case, it is said that the reason that Γ provided is **defeasible**. We say this because it is **defeated** by the presence of a consideration in Δ; we call that consideration a **defeater**.

Before diving into a critique of defeatism, it will be helpful to say a bit more about this view, which is rarely explicitly endorsed and certainly not in such sparse terms. In doing that (the remainder of this section), I'd like to bracket a number of issues related to defeat in order to focus on one important issue. The issues I bracket, however, show that the phenomenon in question is even more variegated and complicated than my main argument acknowledges.[38]

2.3.1.1 What is a defeater

The first issue I'd like to bracket is the question of what a defeater is. If p serves as a justification for S, on the basis of which, S knows q, then it seems natural to think that d is a defeater just in case p together with q does not serve as such a justification. This suggests that d can serve as a further premise (i.e. is the same sort of thing that p is). For example, my belief that Tweety is a bird could provide justification that Tweety can fly, though if I also acquire the belief that Tweety is a penguin we might think that that further *belief* acts as a defeater (the conjunction of the two beliefs does not serve as a justification for the belief

[37]Namely that defeat is not the right or best concept to account for how non-deductive knowledge/justification might be lost in some cases (i.e. that a "no defeater" solution to Gettier misses something). Aarnio (2010), for example, argues that we ought not to be concerned with defeaters, but rather whether knowers are employing *reasonable* strategies for acquiring knowledge. See also her (2014) for similar considerations. Baker-Hytch and Benton (2015) advance a similar argument against defeat though unlike Lasonen-Aarnio also raises problems for internalists who over-leverage the concept.

[38]Throughout this section I was aided by Grundmann (2011); Moretti and Piazza (2018); Sudduth (2017); Koons (2017).

that Tweety can fly).

Of course, defeat is talked about in manifold ways. For example, an externalist may have it that my disposition to think that things which look green are green provides justification for knowing they are green when they are in fact green. But such a disposition will also lead one to think a white object green if it is illuminated by green light. That the object is white and illuminated by green light defeats the justification ordinarily provided by the disposition. But clearly the model of understanding defeat in terms of a kind of inference won't do here: its being a white object illuminated by green light and my disposition to think things which look green are green are very different sorts of things: one is a disposition, the other is a state of affairs or simply some object's being some way.[39]

Even when the defeater is the same sort of thing as that which provides justification, it is not always clear that we should think of it as having the same sort of status. For example,[40] it may be that I have a disposition to always bluff when I have an Ace in my hand in a game of Texas hold 'em. It may also be that I have a disposition to think that a player is bluffing if she checks on the flop and turn and then makes a large bet on the river.[41] Unbeknownst to me, a player has caught on that I tend to bluff when I have an ace. Now, if the flop, turn, and river seem unlikely to support a good hand, and I have made a bet at every opportunity, my disposition may have led my opponent to think I am bluffing (and I am: I have an ace after all). It may also turn out that my opponent does not have a great hand, but has made bottom pair (ordinarily a very mediocre hand). When she raises on the river, I would ordinarily be right to think that she is bluffing (she would be overplaying her hand); however, my disposition to bluff with an ace has given her reason to think that

[39]Some might be tempted to redescribe this case as a case in which one is justified, but the belief is acquired from the reliable disposition is false. If it is a green object illuminated by green light, however, the justification seems accidental—that is, my disposition is not appropriately sensitive to the color of the object since I would have thought a white object in similar circumstances to also be green. A primary motivation behind epistemic externalism is precisely to allow justification to be defeated in these sorts of cases.

[40]Similar examples can be constructed for internalists if we allow ourselves to think of beliefs to act as causes. It is sufficient to allow a belief to figure in an explanation—even when it does not concern the rationality of its possessor.

[41]In Texas hold 'em. There are three rounds of betting during which more information is revealed to all players. If a player has an okay hand, they might bet early to force people without a good hand to fold. If someone only casts a bet on the river, then either they got lucky or they are bluffing. They are often thought to be bluffing because (i) this is their last opportunity to force their opponents to fold, and (ii) it is very unwise to keep playing if one only has a small shot at a decent hand.

bottom pair has a good shot of winning (and it does), so it is no bluff after all. It may seem stretched[42] but my disposition to bluff with an ace acts as a defeater in this case to my disposition to think that a big bet on the river means my opponent is bluffing.

Even though both of the above are dispositions, it seems clear that the role they play in my epistemic circumstance is very different. It seems natural, that is, to think that such a case is different, for example, from a case in which we have two dispositions which conflict with each other (as, for example, my disposition to think birds fly and my disposition to think penguins do not). In order to avoid complications such as this, and to avoid thorny issues that might arise from disagreements about how to understand justification, I will speak of defeaters as **considerations**. A *consideration* is simply something which is relevant to a particular case (regardless of whether we conceive the case internalistically, externalistically, or in some other manner). My hope is that speaking in this broad manner avoids complications that might arise from endorsing a particular kind of account.[43,44]

2.3.1.2 What does a defeater defeat

Second, there's a question concerning what a defeater defeats. I want to bracket one dimension along which this question is typically posed: whether defeat is of knowledge or of justification.[45] Does it generalize over the two or should it be restricted in some further way? These can make substantive differences to accounts of knowledge. If defeat is understood as undermining justification (rather than knowledge), for example, then a no-defeater response

[42] Part of my objection (in the following section) to the notion of defeat is that there doesn't seem a sensible way to exclude this kind of case even though intuitively many philosophers might want to claim that they would want to. The *point* is that my disposition is responsible for the defeat, but only in a deviant manner.

[43] A potential objection is that by broadening what a defeater is in the manner in which I have, I have let so much in that it is unsurprising that I am unable to find a unified phenomenon of defeat. My arguments below do not turn on anything specific regarding defeat as one thing rather than another. That is: they go through regardless of what sort of thing we understand a defeater to be. Since it would be tedious to think of every possible candidate for what a defeater could be, and then to rehearse more or less the same arguments for each of them, I have chosen to formulate things in this way.

[44] Bergmann (2006) makes a distinction between *propositional defeaters* and *mental state defeaters*. The former defeat knowledge (if not justification) in virtue of existing. The latter require awareness on the part of the knower (and defeat justification). This sort of distinction among defeaters does not concern me in this paper.

[45] I have restricted my argument in this section to epistemological defeat, but the question becomes broader if we are talking about defeat more generally. If we turn our attention to defeat of a reason-giving or action-justifying relation (or of actions or reasons), we would face more questions.

to Gettier simply denies the supposed counter-examples.[46] Understanding defeat as defeat of knowledge, by contrast, allows for there to be justified, true belief, which is not yet knowledge. If one understands a consideration as the same sort of thing which may provide justification (as may be more natural for an externalist account of justification), then it might seem more natural to understand such considerations as interacting primarily with justification. If by contrast a consideration is not the same kind of thing as a justifier, then it might be more natural to understand defeat of knowledge. For example, if one has an internalist conception of justification—for example, by understanding a justifier as a belief-like state of a knower—then a consideration could be something external to the agent entirely. My belief that it will rain may be justified on the basis of my beliefs formed in light of a friend's testimony. But if I am unaware that my friend (though reliable in most circumstances) is misled by a bias in this particular circumstance (perhaps he has wagered a lot of money on a sports game in which the team he is backing is known to perform better in inclement weather): such a consideration could defeat knowledge.

But there are other candidates for what a defeater might defeat. For example, perhaps it defeats a consideration (so that one consideration knocks out another). It might defeat a conclusion supported by a consideration. So, for example if p provides reason for q then a consideration could defeat q. Or it might simply defeat the reason-giving relationship itself without passing judgment on q. This relationship will be of central importance to my topic. To get clearer on *these* options, we'll have to say a bit more about how defeat is understood.

2.3.1.3 What is defeat?

Before venturing too far, it's worth pausing to say *what defeat is.* So far I've relied on my reader's familiarity with the concept, which has some ubiquity in analytic philosophy. There aren't terribly many accounts, but here are a few early ones of what counts as defeat.

Chisholm (1966) is arguably the first use of the concept of defeat within epistemology. He defines defeat as follows:[47]

[46] Externalist theories of justification are motivated in part by an attempt to do this.

[47] It isn't important for my purposes, but in case anyone is curious Chisholm defines "tending to make evident" as:

The relevant concept of *defeat* is this:

D3 d [defeater] defeats e's [evidence] tendency to make h [hypothesis] evident [$=_{df.}$] e tends to make h evident; and $d\&e$ does not tend to make h evident.

It's important to notice two trends in Chisholm's definition. First, the left hand side suggests that defeat is a relationship between some further consideration and a particular justificatory relationship (namely the relationship that e has to h). The second trend is found by examining the right hand side of the definition. Namely, that d doesn't so much as defeat the relation that e has to h in its presence—the right hand side seems to avoid taking a stand here—rather, its presence (taken together with the evidence e) undermines the status itself. That is, we can make a distinction between a justificatory or reason-giving relationship that a consideration (or collection of considerations) has towards some particular proposition. And the status of all considerations on an occasion actually supporting (or not) some particular proposition. The left hand sides suggest the first *relational* notion. The right hand side suggest a functional notion (we get an extension of cases without an explanation of the mechanism).

A similar sort of ambiguity runs through a large number of accounts in the past 50 years. The issue is that epistemologists seem to *want* to understand defeat in terms of the undermining of a relation of justification, but set out to do this by understanding it in terms of the undermining of the status of justification that all of one's considerations have. If it turns out these are the same, that wouldn't be problematic (but as I hope to show that's not the case). Here, are two more accounts, from that era:

Lehrer and Paxson Jr (1969), who are borrowing explicitly from Chisholm, and whose paper is the epitome of the no-defeater-response-to-Gettier, write (bolding added by me for emphasis):

> We propose the following definition of defeasibility: if p completely justifies S in believing that h, then **this justification is defeated** by q if and only if (i) q is true, (ii) the conjunction of p and q does not completely justify S in believing that h, (iii) S is completely

---

D2 e tends to make h evident $=_{df.}$ e is necessarily such that, if it is evident for someone x, then there is a w such that (i) w is evident for x, (ii) $w\&e$ cannot be anyone's total evidence, and (iii) $w\&e\&h$ can be someone's total evidence.

In short: you can't add e to your "total evidence" without bringing along *something*, and one such something happens to be h (though there may be other such somethings).

justified in believing q to be false, and (iv) if c is a logical consequence of q such that the conjunction of c and p does not completely justify S in believing that h, then S is completely justified in believing c to be false.

Klein (1976) comes close to avoiding this ambiguity, but drops the ball at the end of his definition:

> S knows that p if and only if the traditional conditions hold [justified true belief] and there does not exist a true proposition, d, such that if S were to add d to whatever justified p for S, P would no longer be justified for S. We can rephrase the last condition as: There does not exist a true proposition, d, such that $(d \cdot e)$ fails to justify p for S (where e is the evidence S has for p such that e justified p for S). I will use 'xJ_sy' to mean 'x justifies y for S' and '$x\overline{J}_sy$' to mean x fails to justify y for S'. If xJ_sy, and d is true and $(d \cdot x)\overline{J}_sy$, d **is a defeater of the justification** of y by x for S.

Pollock (1974) writing at around the same time noticed a similar ambiguity and introduced a distinction between "rebutting" and "undercutting" defeat. Here are Pollock's definitions:[48,49]

[48]These formulations are lifted from Pollock and Cruz (1999). Sturgeon (2014); Melis (2014) have argued that undercutting defeat can be assimilated into rebutting defeat by understanding the former as a case of the latter concerning a higher order belief about evidential relations. Casullo (2018) undercuts this argument. For reasons that will become clear by the end of the paper, I am skeptical of so-called "higher-order beliefs". I don't think it's impossible to form beliefs about one's evidence. But the risk is misunderstanding evidential relations in such a way that they require higher order beliefs (e.g. to conclude q from p requires a belief of the form "if p, then q").

Pryor (2013) challenges whether Pollock's distinction is exclusive.

[49]Here are the exact formulations employed by Pollock. Since I am trying to remain neutral on issues concerning internalism/externalism, I found it easier to ignore his bifurcated definitions:

Reason A state M of a person S is a *reason* for S to believe Q if and only if it is logically possible for S to become justified in believing Q by believing it on the basis of being in the state M.[50]

Defeat If M is a reason for S to believe Q, a state $M*$ is a *defeater* for this reason if and only if the combined state consisting of being in both the state M and the state $M*$ at the same time is not a reason for S to believe Q.[51]

Next, Pollock introduces a distinction (important in the literature) that comes close to capturing the distinction I've been gesturing at. This is the distinction between "rebutting" and "undercutting" defeat:

Rebutting Defeater If M is a defeasible reason for S to believe Q, $M*$ is a *rebutting defeater* for this reason if and only if $M*$ is a defeater (for M as a reason for S to believe Q) and $M*$ is a reason for S to believe $\neg Q$.

Undercutting Defeater (for doxastic state) If believing P is a defeasible reason for S to believe Q, $M*$ is an *undercutting defeater* for this reason if and only if $M*$ is a defeater (for believing P as a reason for S to believe Q) and $M*$ is a reason for S to doubt or deny that P would not be true unless Q were true.

Undercutting Defeater (for non-doxastic state) If M is a nondoxastic state that is a defeasible reason for S to believe Q, $M*$ is an *undercutting defeater* for this reason if and only if $M*$ is a defeater (for M as a reason for S to believe Q) and $M*$ is a reason for S to doubt or deny that he or she would not be in state M unless Q were true.

Rebutting defeaters defeat the truth of the conclusion supported by the evidence (without necessarily undermining the evidence itself). That is, if p provides reason for q, then d is a rebutting defeat if it defeats q. Often this is cashed out in terms of the idea that d provides (stronger) reason for $\neg q$.

Undercutting defeaters undermine the way in which the evidence supports the conclusion. NB: typically they are understood as *thereby* defeating the conclusion, but this isn't necessarily so.[52]

Bergmann (2006, p.159, n.12) introduces a notion of "reason-defeating defeater" to put alongside Pollock's distinction above:

Reason-Defeating Defeater defeats the truth of the claim which provides reason (without thereby defeating the reason giving relation itself). That is, if p gives reason for q, then d is a reason-defeating defeater if d defeats p. Nevertheless, it is still possible that p *would* give reason for q.[53]

Grundmann (2011, p. 158) provides the following useful disambiguation of these three notions:

> "You travel through a particular country district and up to now you have only seen animals that you believe to be brown cows. That inductively justifies your belief that the cows belonging to your friend Jim, who lives in this district, are also brown. The justification of your belief can now be defeated in the following three different ways: 1. While you visit Jim at his farm it turns out that he does not possess brown cows, but only black and white ones (rebutting defeater). 2. On your further trips through this district you realize that your observations so far have not been representative. Many farmers in the district also possess black and white cows. For this reason, your observations no longer inductively justify the belief that Jim possesses brown cows (undercutting defeater). 3. You find out that your original perceptual beliefs were false. You haven't seen any brown cows, only cleverly disguised black and white ones. Someone has played a joke on you. What you took for a reason turned out to be false (reason-defeating defeater)."

[52] A recent paper by Muñoz (2019) argues for a separate phenomenon, *disqualification*, which defeats a reason giving relation without defeating the thing which is supported by the reason (i.e. the conclusion). A paradigmatic example is a case where previously q was supported on the basis of testimony, but then we receive direct evidence of q. The direct evidence disqualifies the testimony without defeating q (since it also provides evidence for q). Part of Muñoz's argument relies on the idea that undercutting defeat must defeat not only the reason-giving relation, but thereby the conclusion. I don't see why we have to understand things in this way. Regardless, "disqualification" is an interesting phenomenon to lay alongside all of these other ones.

[53] This claim itself requires a lot of nuance to cash out. What does this "would" amount to? Surely anything that provided support for p would defeat this defeater? So, how can we say that absent that support the reason-defeating defeater leaves the reason relation in tact?

There is a huge literature on the various sorts of defeaters and defeat like phenomena there could be. I have so far surveyed a number of distinctions that are important for present purposes.[54] In this section I rehearsed a number of different accounts of defeat as well as a number of distinctions among kinds of defeaters. The result is that the landscape is quite complicated. Defeaters can take the form of all sorts of considerations and can defeat propositions, reason-giving relations, reasons (but not the corresponding relations), and so on. There has been a tension lurking among all of these notions. Often we define a defeater at least partially functionally, i.e. in terms of a reason-giving relation wherein the supported conclusion is no longer supported. Next, we specify different species of belief by saying *what further thing* they do (apart from rendering the conclusion no longer supported). Or, to be precise, the definition is functional, but is filled in with something which is relational (the defeater is related to the evidence in a particular way, namely that of defeating it).

I will refer to these as **functional** and **relational** notions of defeat, but I should be clearer on what is really at issue. The functional notion of defeat involves what we might call an entire circumstance or an entire evidential circumstance. If Γ is our evidence (or, at least, the set of relevant considerations), and Γ provides reason for p, then a defeater may be

[54] I hope that my coverage, through footnotes is more or less exhaustive, but I'm sure I've missed something. A few other examples of note. First is the notion of a "no-reason defeater" from Bergmann (1997). These are defeaters of beliefs (or propositions) for which has no evidence/reason. The intent is to capture the thought that we may acquire information which defeats an epistemologically basic belief (for example, a perceptual belief—provided one allows for basic beliefs).

- Unlike reason-defeaters they don't defeat reason for an inferential belief (i.e. a belief which provides reason for something further)
- Unlike undercutting beliefs they don't sever the justificatory link that the belief they defeat *would* provide.
- Unlike rebutting defeat they don't provide any reason for thinking the defeated belief is false

Finally, Lackey (2008) introduces two further sorts of defeaters. These concern the status of a defeater. I have used the locution, "consideration," but must considerations be available as evidence?

Doxastic Defeaters are **beliefs** which speak against a belief or the reliability of the evidence in favor of that belief. Notably doxastic defeaters require no evidence themselves to defeat one's beliefs. Nonetheless, it seems intuitive that possessing them should defeat one's knowledge/justification.

Normative defeaters are beliefs that one *ought to acquire* which would undermine/defeat one's belief. They are similar to propositional defeaters (insofar as they lie outside what one has the right kind of access to—though this needn't be the case if for example one has some kind of bias against entertaining evidence that she in fact has—but they are beliefs one *should* have if they did their due diligence or didn't have certain biases, etc.).

Normative defeaters are particularly interesting and have spawned some discussion (Goldberg, 2016, 2017; Benton, 2016).

defined functionally in terms of what happens when we add an additional consideration to Γ. That is, d is a defeater in case when $\Gamma \cup \{d\}$ is our set of considerations, p is not so provided for.[55]

By contrast, when we say that Γ provides for p, what we mean is that Γ provides reasons for p. Perhaps q in the context of background information k (which are both contained in Γ) provide this a reason for p. Further: no other information in Γ provides a stronger reason against p (or perhaps any reasons against p). Hence Γ provides reason for p. These notions are relational. They concern a **particular relationship** between considerations and not the outcome of all of those considerations. A defeater may defeat the particular reason-giving relationship that q against background k has with p.[56]

[55] I am here relying on a distinction between contributory reasons and overall reason. See Dancy (2004b). Dancy plays a large role in the arguments of this section. I return to Dancy in the fourth chapter of this dissertation (where I also spell out this distinction in more detail).

[56] Kvanvig (2007) notices a similar trend and a similar problem. I believe that Kvanvig advances a similar argument against "functional" accounts as I do (in the sections following this). I'll quote at length an illuminating metaphor Kvanvig uses for these two approaches. His "frontdoor" approach is my relational notion; his "backdoor" approach is my functional notion. The difference here is whether we understand changes to our entire evidential circumstance as prior to a characterization of defeat (defeat is characterized in terms of such changes), or whether the characterization of defeat is prior (and used to characterize changes to that whole):

> "Here, I want to explore two fundamentally different approaches to the concept of defeat, and argue that only one of them has any hope of success. One theory begins with propositional relationships, only by implication describing what happens in the context of a noetic system. Such a theory places information about defeat up front, not informing us of how the defeat relationships play out in the context of actual belief, at least not initially. The other theory takes a back door to the concept of defeat, assuming a context of actual belief and an entire noetic system, and describing defeat in terms of what sort of doxastic and noetic responses would be appropriate to the addition of particular pieces of information. Where the house is the noetic structure itself, the front-door approach characterizes the concept of defeat in terms of the propositional contents a belief might have, thus characterizing defeat at the front door. It presumes that once let into the house, some changes will be required, but the characterization of defeat is logically prior to any account of such changes. The backdoor approach characterizes defeat in terms of what leaves the house, in terms of beliefs that exit the noetic system in response to intrusions into the system, in terms of what the staff of a well-run household kicks out the back door for making a mess of things. The best-developed example of a backdoor theory is Alvin Plantinga's, and here I will argue that his theory and approaches like it will be unable to explicate accurately the concept of epistemic defeat. I will argue that a front-door approach is needed, rather than a backdoor approach." (Kvanvig, 2007, p. 108)

2.3.2 Defeating Defeat: The Feat To be Accomplished

In the close of the previous section, I argued that definitions of defeat seem to vacillate between two notions: a relational notion of defeat whereby a defeater defeats some particular thing, and a functional notion whereby a defeater is understood in terms of the effect it has on our overall evidential circumstance (and not necessarily particular considerations or relationships between considerations). It is worth observing that these two notions are indeed distinct. They aren't just notionally distinct, but extensionally distinct as well. The following two examples demonstrate this.

Party Guests Susan is an excellent party guest. If Susan attends a party, a person has good reason/justification for thinking that the party will go well. Bradley, by contrast, is a terrible party guest (loud, toxic, rude, etc.). If Bradley attends a party, a person has good reason/justification for thinking that the party will go poorly. As it turns out, lovely though Susan may be, she is not nearly excellent enough to counter Bradley's toxicity. If both Susan and Bradley attend a party a person does *not* have reason/justification for believing that the party will go well.

This is an example that fulfills the functional definition of defeat: the conjunction of the defeater and the original belief are insufficient to justify the original conclusion. Nevertheless, the original justification still stands: it is still the case in such circumstances that Susan's presence provides something in favor of thinking the party will go well. That is: even while Bradley is there, we can still think the party is better now that Susan is there. This relationship holds even if it is not decisive.

We can compare this with a case in which Bradley's presence is known to have an incendiary effect on Susan: when Bradley shows up, Susan is irritable and argumentative. She and Bradley often get into screaming matches. In this case, Bradley's presence defeats the relation itself. In PARTY GUESTS Bradley's presence does not work in this way: Bradley is more toxic than Susan is pleasant. We can even make the situation more interesting in the following way: we know that Susan is *such* a good party guest that when a disruptive or toxic person is present, she tries even harder to make the party go well. In such a circumstance Bradley's presence actually *strengthens* the case that Susan's presence makes for the party going well (whatever the case is for Susan's presence making parties go well, we suppose that she's even more effective when someone like Bradley is there), nevertheless Bradley is such a toxic person that even a *strengthened* Susan fails to provide a sufficient/reason justification for thinking the party goes well.

I believe the lesson from this case is clear: Susan's presence is a reason/justification for thinking a party will go well; Bradley's presence is a consideration which defeats this. It defeats it not by undermining the reasons that Susan's presence provides, but rather by making a stronger case against the party going well.[57]

We can also find examples of defeat that fall under the relational notion, but not the functional notion. The structure of the following example is that a defeater defeats a particular relation of reason/justification but nevertheless itself provides a reason for the very same belief:

Car Pool I'm curious whether Trevor will be at my party on Saturday. My friend, Mark, reports that he is giving Trevor a ride there. Mark's testimony provides a reason/justification for believing/knowing that Trevor will be at my party. As the party is beginning, I go outside to greet guests. When I'm outside I see a different friend, Felicia, driving towards my house with Trevor in the front passenger seat. This new information defeats the original reason/justification that Mark's testimony gave for believing/knowing that Trevor would attend. Nevertheless, the information *also* provides a reason/justification for believing/knowing that Trevor will be at my party.

In the above example the particular reason/justification relation is defeated, but a similar relation is also generated. For this reason, the belief/knowledge is not defeated (though a particular relation is).[58]

Clearly these two notions are distinct. So we cannot rely on the functional notion to explicate the relational notion. It may be that the functional notion captures all and only those relational notions of defeat which end up disturbing what is supported by all of the

[57] Within the functional notion, the distinction I am invoking here is between undercutting and rebutting defeaters.

[58] Muñoz (2019) argues that we should understand this sort of case in terms of a notion he calls *disqualification*, which he lays alongside undercutting and rebutting defeat. For me: disqualification is a certain kind of undercutting belief, namely one in which the entire circumstance ends up supporting the same conclusion.

relevant considerations. In this case we should treat the definitions of defeat not as descriptive definitions—which seems to require at least extensional adequacy—but rather as explications of defeat: they seek to take some set of things isolated by the functional notion and provide an explication of that class. It might not be extensionally adequate, but it explicates what defeat is about. In order for this suggestion to work, we must have an independent grasp of the relational notion. In §2.3.2.1, I show this to be untenable. Following this, I examine a few conceptions of the functional notion. My aim is start with the notion that appears in the literature and show that we must weaken it until what we have can no longer plausibly be called defeat. I cannot claim this is a knockdown argument against defeatism, but I hope to show a way of viewing the phenomena that is more apt.

2.3.2.1 A Relational Notion of Defeat

Here I argue against what I call a relational notion of defeat. That is: while it might appear at first glance to be a unified phenomenon, it is actually just a collection of a large number of divergent things working in radically different ways. My suspicion is that focus on very simple cases has led epistemologists to think that the notion is unproblematic. Whatever the reason for its current prominence, I hope my conclusion will become clear: such a notion is chimeric.

To start, I'll need to set a bar for adequacy: an account—however vague, promissory, or Frankenstein—of how it is that a consideration's presence can undermine a relation of justification. What sort of effect does it have on one's evidence or reason/justification such that that reason/justification no longer stands in such a justificatory relation.

I'll start with a rather obvious and intuitively broad thought: it works like a light switch, it simply turns off or *disables* the way in which something provides reason or justification:[59]

Disabling Defeat: d is a disabling defeater for p's being a reason for q if d is a consideration whose presence makes it such that the relation no longer holds.[60]

[59]This section appropriates heavily from (Dancy, 2004b). In particular, the three most prominent candidates of defeaters in this section—disabling, weakening, and contradicting—required very little innovation to construct from Dancy's materials.

[60]This notion is equivalent to the notion of "undercutting defeat" if such a notion doesn't require the entire reason relation to change. Otherwise, this notion is is distinct.

For example: that an object appears red to me provides a reason/justification for believing/knowing it to be red. That the object is illuminated solely by red light, however, is a consideration which defeats this justification relation: that the object appears red no longer provides reason/justification for thinking it red.

Unfortunately, this is not the only way in which a consideration can act to undermine a relation of justification, and importantly not the only way it can do this in what we would otherwise consider a case of defeat. For example, a consideration may simply sufficiently weaken (though not negate) a reason/justification relation:

Weakening Defeat: d is a weakening defeater for p's begin a reason for q if the presence of d sufficiently *weakens* the relation such that p is no longer a sufficient reason/justification for q.

For example: Suppose that Ben sees a structure that appears to be a barn. This gives Ben a reason/justification for thinking it is a barn. Now suppose that Ben talks to an epistemologist who warns him that there are Barn facades. Given this new information, Ben should be *less certain* that what he sees is a barn, but it isn't sufficient to defeat his evidence. Of course, we can incrementally weaken the original evidence (and thereby make Ben less certain). Suppose Ben learns that there are n barn facades within x miles. Increasing n or decreasing x sufficiently will eventually make Ben's original reason: seeing a structure that appears to be a barn, sufficiently weak that it no longer provides a reason/justification for believing/knowing it to be a barn. The important point here is that the appearance of a barn doesn't *cease* to provide justification at all—as was the case with our red object in red light (the information no longer provides justification)—rather, the prominence of barn facades weakens the degree to which the evidence provides justification, and this can happen to a high enough degree that an agent lacks sufficient reason/justification.[61]

[61] In addition to these two types of defeaters, a reader of Dancy might extend this reasoning to a third kind:

Contradicting Defeat: d is a contradicting defeater for p's being a reason for q if the presence of d provides a sufficiently strong reason/justification against q that a person, ought not to believe q given d: that is: $p\&d$ do *not* provide reason/justification for q (in some cases they might even provide reason/justification for $\neg q$ where the contradicting defeater is sufficiently strong).
Example: Party Guests from above.

This isn't a clean distinction; sometimes it's hard to categorize things. For example: suppose that an object appears green. This is a reason/justification for believing/knowing it to be green. If we learn, however, that

Now, it seems clear that disabling and weakening defeat are distinct and work in distinct ways. A sufficiently clever philosopher might intuit some common mechanism. Some light switches are dimmers. So you can flick a light switch on/off, or you can dim in to the point that it no longer detectably emits light; alternatively: we might simply adopt a disjunctive account of the two. Regardless, we can come up with a number of other ways in which the presence of some consideration works to undermine a justificatory relation, which nevertheless fails to satisfy either of the above. I'll quickly walk through two examples and then try to illuminate some lessons for how we can construct more:[62]

Weakener Enabler: d is a weakener enabler just in case d is a consideration whose presence renders some further consideration d_2 a weakener sufficiently strong to undermine a justification relation, but such that d_2 need not act as a weakener in the absence of d.

Example: A modification of PARTY GUESTS has this feature. Suppose that Mary is ordinary a great party guest; likewise, suppose that Bradley is ordinarily an okay party guest—for simplicity suppose that Bradley's presence does not bear on whether a party will go well or poorly, but Mary's presence provides reason/justification for thinking it will go well. We can suppose further that Mary and Bradley get along fine... unless there is alcohol involved. Now, it isn't that alcohol makes Bradley an atrocious party guest. In fact, alcohol often improves Bradley's disposition. Nevertheless, Bradley's mannerisms under the influence of alcohol have such an effect on Mary that she is no longer the pleasant party guest she ordinarily is. In such a

it is illuminated solely by green light, then this is a disabling defeater: its appearing green no longer provides a reason/justification.

However, if instead the object is illuminated solely by yellow and blue light, this actually provides a reason *against* thinking it green (a green object in such a circumstance would appear black: it would reflect no light, a white object would appear green). It's unclear how to categorize this kind of defeat. It's not a contradicting defeater because it only provides a justification for the opposite conclusion in conjunction with the original reason (not on its own). It is probably a weakening defeater. If that's right, then the lesson is that weakening is not always a matter of degree (sometimes it's on/off). The point is that the object's appearing green becomes a reason *against* thinking it green (i.e. it does more than cease to be a reason in the presence of the consideration).

It is also worth noting in this connection that Pollock's distinction above—between rebutting and undercutting defeat—doesn't map cleanly onto anything I've mentioned, though is perhaps analogous to contradicting and disabling defeat, respectively, in spirit.

[62] I'm careful to choose the word "lessons" here. The whole point is that the whole matter is such a big hurly-burly, that we can't hope to get a grasp on anything like principles for outlining a class like "defeater" (in the relational sense).

circumstance the availability of alcohol enables Bradley's mannerisms, which serve to sufficiently weaken the case that Mary's presence makes for the party going well. In other words: the availability of alcohol is a *weakener enabler*.

To continue the analogy from previously: turning on a light in another part of the house need not sufficiently dim this light; however, if we are using a sufficient amount of power, then turning on that light *will* cause the light in this room to sufficiently dim. In this case the fact that other sources in the house are already drawing a lot of power (perhaps together with how the house is wired) is a weakener enabler of something else: namely that turning on a second light causes the original light to sufficiently dim.

Weakener Strengthener: d is a weakener strengthener just in case d is a consideration whose presence renders some further consideration d_2 a weakener, but such that d_2 is only sufficiently strong (as a weakener) in the presence of d.

Example: I have a day off of work. Suppose this gives me reason to go on a hike. So my day off provides a reason for the hike. I look at the weather report and see that based on barometric readings, there is a modest chance of precipitation. Ordinarily, rainy hikes are fine. They provide a unique challenge. But the day off provides support for a hike because a hike would be relaxing. What this means is that rain doesn't count against hiking, but it *does* count against the way in which a day off provides support for taking a hike. Nevertheless, the chance of precipitation is pretty modest. It weakens the case for taking a hike, but not enough that the day off no longer counts in favor of the hike. Suppose that in light of further evidence (satellite images) the chance of precipitation is thought to be quite a bit higher. Then this strengthens the way in which the chance of rain weakened the case that my day off made for taking a hike.[63]

These are of course considerations which work in a completely different manner, but nevertheless satisfy what I was called the "relational notion". In addition to these, it is possible for a consideration to enable an disabler (though not itself serve as a disabler); or strengthen a disabler (without itself serve as a disabler); or weaken an enabler (without itself serving as

[63] I won't even try to make an electrical analogy here.

a disabler); etc.. There are several further ways that things can get complicated.

- The previous considerations we might call "2^{nd} order" in virtue of the fact that they bore on other considerations in a way that led to a reason-giving relation being severed. But we could have 3^{rd} order or arbitrarily higher order considerations.
- In addition considerations can work to change the polarity of other considerations (not just strengthen/enable), but change a weakener into a strengthener.
- It is likewise not clear that these categories are exhaustive (so that we have a way of recursively enumerating all such considerations.[64]
- Finally, that a consideration d is a defeater of p's being a reason for q may change in unpredictable ways given further considerations. It may be that against background information k, d is understood in this way, but further information changes things. It might not be the case that this further information is a consideration which bears on the way that d bears on p's being a reason for q (but just changes the structure of the situation sufficiently that it no longer so bears).

My claim is that we cannot in general say what it is for a consideration d to be a defeater of p's being a reason for q. What we in fact find is a grab bag of different ways in which considerations interact with each other. While there are systematic things we can say about how these things interact, we cannot isolate any phenomenon within that messy web of interaction that corresponds to defeat. We in general cannot account for the way in which one piece of information might bear on another.

2.3.2.2 A Functional Notion of Defeat

In light of the previous, we might think that even though defeat is not a unified phenomenon, we might still just take whatever the functional notion gave us and call that defeat. We might just think that the differences between what the functional notion delivers and what we intuitively took defeat to be can be explained away via something like Carnapian explication. This idea has something to it, insofar as the extensional differences I invoked

[64] Hart's original point about defeat—and why he felt that there was a lacuna in legal conceptual space—was that the phenomenon in question (in contradistinction to e.g. a negative consideration or desideratum) is open-textured.

above don't seem that problematic for the phenomenon. One might think that, although cases like Party Guests aren't strictly included in the exemplary cases of our pre-theoretic conception, they nevertheless belong in a class with such cases insofar as they involve some further consideration undermining or defeating *something* (even if not the justification relation). Likewise, one might be fine simply dismissing as errant cases like Car Pool on the grounds that while they do involve the relational notion, they aren't in the target notion: namely cases in which the overall verdict that all of our considerations yield is actually defeated. So: even though the exemplary cases of defeat seem to involve a relational notion of it, we might nonetheless think that the *real idea* we were after is captured (at least extensionally) by a functional characterization. That it isn't exactly the same isn't a problem since it captures the central cases we care about and in a number of cases of like-phenomena.

Unfortunately, the functional definition yields some strange verdicts about what is defeating what (or what counts as a defeater). While a traditional epistemologist might be prepared to dismiss Car Pool and include cases like Party Guests on the grounds that the former isn't actually in the target notion, and that the latter is a like phenomenon, the sorts of cases I have in mind in this section aren't so easy to assimilate. In particular, I have in mind cases in which the functional account yields the incorrect verdict. Basically cases where the traditional account is led to say "defeat happened", but is in an awkward position to say e.g. what the defeater is.

To be precise, we are supposing that we understand defeat as:

Functional Account: Suppose Γ provides reason for p. Suppose that $\Gamma \cup \{d\}$ does not provide reason for p. Then we call d a defeater. (Since the original justification is defeated)

I think you can most clearly see this in the following kinds of examples. In the examples in question p serves as a reason/justification for q. d is not a defeater: $p\&d$ also serve as a reason—we can even imagine that d contributes positively here, though it's not necessary to set up the case in this way[65]—nevertheless, e turns d into a defeater for the original relation

[65]It works best if p is insufficient but together with d is sufficient, though e (i) makes p sufficient while also (ii) rendering d a defeater. According to a functional account e is the defeater, but e is working more like an disabler enabler (to transpose the vocabulary of the previous section).

(between p and q). According to the functional definition we should understand this as e defeating $p\&d$.[66,67]

Cloth in the Window: You and a friend have agreed that a green cloth in a shop window is to be used as an important signal for q. The conjunction of your knowledge about this plan (p) and the appearance of a green cloth in a particular shop window (d) gives you a reason/justification for believing/knowing that q. You pass by the shop everyday seeing various colors of cloth (some days no cloth). Eventually the day arrives where you see what appears to be a green cloth (d). Unfortunately your friend has limited control over the shop display. She may place a cloth in the window, but doesn't control e.g. the lighting used. Now, suppose that you learn that the shop keep only uses exclusively yellow and blue light (e) (so the yellow cloth you saw was probably yellow; likewise with the blue cloth), but this actually means that its appearing green is a reason against its being green (a green cloth reflects neither yellow nor blue light; but a white cloth would reflect both and appear green).

According to the functional notion the presence of exclusively yellow and blue light defeats the status of justification for the claim p (because $p\&d$ clearly give reason for q but $p\&d\&e$ do not). But this seems to be a weird way to describe the case. For surely, what defeats that status is its appearing green. Under such considerations, that it appears green undermines the status of justification that our belief had.[68]

[66] Couldn't we understand $d\&e$ as defeating p? We could, some accounts allow this, though note that Chisholm's actually rules out understanding it in this way.

[67] Put another way: e turns d from a partial supporter for q into a defeater.

[68] **?** notices a similar phenomenon and uses it to leverage an objection against Pollock's account of (undercutting) defeat. Namely, that if E is evidence for H and D is a defeater of this (i.e. $E\&D$ do not support H), then provided that D on its own supports H, we have that E is a defeater of D's support of H. But we can construct cases where this seems rather implausible as a description. Either we must deny that defeat actually occurs (which spells trouble for the functional definition) or claim that something is a defeater which is clearly not one (also spells trouble for the functional definition). This is very close to what I claimed to show. Here is Chandler's example:

> "outside the door to Sam's flat is a switch for the light on the landing of the floor below. Flipping the switch (E_S) typically causes the light to go on (H_S): E_S is therefore a *prima facie* reason to believe H_S. Of course, in the event of a power cut (D_S), E_S loses this probative force since, under such circumstances, whether or not the light is on is causally independent of the position of the switch. D_S, we would intuitively like to say, is a defeater for E_S's support for H_S. Now, in order to comply with local regulations, Sam's landlord has installed a backup power system that is activated in the event that the main system

One might be tempted to object that it *is* the new belief e concerning the lighting that is the defeater. In fact such a defeater has a name, it is a *reason-defeating defeater*.[69] Such a defeater acts by defeating the truth of a belief which, were the belief true, would provide justification for the conclusion. In this case it acts to defeat the belief that the cloth is green. I think this misses what is important about this case. In a typical case of this kind, we have it that e.g. the belief that an object is red can be defeated by the belief that it is illuminated solely by red light. It is undermined because the red appearance could be the result of it actually being red or being white. That is: the belief concerning the lighting makes it such that the red appearance no longer provides useful information concerning whether the object is red. But this is not how the Cloth in the Window example works. In this case, that it is illuminated by yellow and blue light doesn't make it such that its green appearance fails to provide useful information concerning whether the object is green. Rather: that it appears green tells us that it definitely *not* green—for it if were it would appear black. The purported defeater—or *reason-defeating defeater*—is the cloth's green appearance.

The point of the example is this: the functional notion of defeat doesn't pick out a defeater. Adding a new consideration may make it such that a set of considerations no longer provides reason for some belief; but that doesn't mean that the new consideration is a defeater. It may be that some other consideration becomes a defeater in light of the new consideration, as in Cloth in the Window.

> fails, automatically powering the light in the staircase to prevent tenants and their visitors from stumbling in the dark. So, just like E_S, D_S provides a reason to believe H_S. But it is worth noting an important asymmetry here: while D_S is a defeater for E_S's support for H_S, E_S is clearly not, from a pre-theoretical viewpoint, a defeater for D_S's support for the same proposition: the evidential connection that obtains between the occurrence of a power cut and the staircase light's being on is entirely unaffected by the position of the switch." (**?**, p. 50)

To summarize: clearly a power outage defeats the support that flipping the switch gives to the light's being on. The setup also seems to force us to say that flipping the switch defeats the support that the power outage gives to the light's being on, but this seems to be the wrong verdict.

[69] See above.

2.3.2.3 An Existential, Functional Notion of Defeat

If my arguments have been convincing so far, then I think the notion of defeat stands on the precipice of defeat. Nevertheless a defeatist might contend that although it is difficult to say in advance what consideration is a defeater, that it is defeat which accounts for the target phenomenon is clear; after all, in all the cases so far considered it was the presence of *some* consideration, which explained why some set of considerations provided reason, but an expanded set of considerations did not. In order to defeat this last stalwart, I take aim at the following notion of defeat:

Definition 2.3.2 (Existential, Functional Notion of Defeat)**.** Suppose that a set of considerations Γ provides reason for p. Now suppose that Δ ($\Delta \supsetneq \Gamma$) does not provide reason for p. This is because Δ contains *some* defeater d that defeats the reason Γ provided.

The case I have in mind is one in which a set of considerations provide a reason for p. Acquiring new information renders it such that we no longer have reason for p, but not because the new information (or any of our old information) speaks against p; rather, because of how all the information is structured, it just makes it such that p is no longer supported. The example is specifically about reasons for action, but it is easy to reformulate such a case epistemically.[70]

[70] I am indebted in this section to Titelbaum (2019). Titelbaum presents an argument against the "weighing conception of reason", which he cashes out as at least committed to the following sort of claim: if a set of considerations Γ is said to provide reason for p, then at least one of the considerations in Γ must provide a reason for p; that is: something in Γ must speak in favor of p. He presents the following example which I found helpful for construction my own example here:

Dinner Wine: You have been invited to a dinner part and are tasked with bringing the wine. You must choose among a red, a white, or a rosé. The hosts will serve either steak or a delicate fish. Depending upon what they serve the pairing of your wine choice with their choice of entrée yields the following utilities:

utilities	steak	delicate fish
red	5	-10
white	-10	5
rosé	-1	-1

It seems we have reason to bring rosé on decision-theoretic grounds, but nowhere do we locate a consideration that speaks for rosé. Officially outcomes (represented by the numbers) provide reasons for (if positive) or against (if negative) a particular act (represented by the rows: bringing red/white/rosé). In fact we find plenty of considerations that speak against bringing rosé to the dinner party. Nevertheless, because of how the various considerations hang together rosé is the thing to bring. My example is meant to play on something similar: that while the various considerations of a case coalesce in a defeasible manner, none are defeaters. Rather, it is how they hang together—how they are structured—that explains the apparent defeat.

Weekend Excursion: Ted is planning a weekend excursion. Giving his location he has three options, he can drive to the big city, drive to the mountains, or drive to the beach. There are many things which speak in favor of each of them, but the big uncertainty concerns the weather. For example, if it snows, the city will become congested and more difficult to navigate. If it rains, the mountain becomes muddy and Ted will be forced to stay indoors. And if there's any precipitation (rain or snow), then we can still relax at the beach house, but it's not nearly as nice. Based on these factors Ted sees his utilities as depicted in the utility table in Table 1: Because Ted takes a large number of weekend

Table 1: Utility Table: Weekend Excursion

Action	Rain	Snow	Sun
Big City	80	50	90
Mountain	20	90	90
Beach	10	10	100

excursions, he makes his decision according to the expected utility each action has (Ted figures that strategies like maximin are great for one-off decisions, but the rare unpleasant weekend is worth it if Ted can also enjoy amazing weekends when it is sunny and more often than not have a good weekend regardless). On Monday Ted checks the forecast for the upcoming weekend and finds that there is a very high chance of precipitation and most likely snow. They are calling for a 35% chance of rain, a 60% chance of snow and a 5% chance of sunny weather. Based on this forecast going to the mountains is the thing to do. However, suppose that later in the week, on Thursday, Ted learns that a change in weather has altered the probabilities. The forecasters are saying there's a good chance the weather system will dissipate before it reaches Ted and sunny weather is to be expected. They now forecast a 35% chance of rain, a 60% chance of snow and a 5% chance of sunny weather (see Table 2. Based on this new forecast Ted concludes that he shouldn't go to the mountains, but instead should go to the city (expected utilities depicted in Table 3). If it rains that will be bad, but Ted knows the risk is worth it, he has done the calculation.

Table 2: Weekend Excursion: Monday and Thursday Probabilities

Forecasts	Rain	Snow	Sun
Monday	35%	60%	5%
Thursday	20%	20%	60%

Table 3: Weekend Excursion: Expected Utilities

Expected Utilities	On Monday	On Thursday
Big City	60	**80**
Mountains	**61**	76
Beach	14.5	64

According to the defeatist something among our considerations is acting as a defeater for the reason Ted had previously to go to the mountains. We might be tempted to say it is the new forecast, but the new forecast only gives Ted *more* reason to go to the mountains; it only *strengthens* the case for going to the mountains. Still, because of how all of the considerations hang together, how they are structured, what we had reason for before, we no longer have reason for with our expanded set of considerations. But this isn't because anything (among the old or new considerations) acts as a defeater. It's just the way that the information is structured.[71,72,73]

2.3.2.4 An Un(der)specified Notion of Defeat?

All that seems left then is what I'll call a un(der)specified notion of defeat since it does not specify what counts as the defeater. It simply gives conditions for when defeat happens.

Definition 2.3.3 (Un(der)specified Notion of Defeat)**.** Suppose that a set of considerations Γ provides reason for p. Now suppose that Δ ($\Delta \supsetneq \Gamma$) does not provide reason for p. This is because *there is defeat*. We therefore call the reason that Γ originally provided *defeasible*.

Such an account fails to count as defeatism on my view, for it falls short of the claim that there is some consideration—a *defeater*—which does the work of defeating. Instead it simply acknowledges that there is a phenomenon in which the addition of further considerations may have a negative effect on which conclusions are supported by reason. This un(der)specified notion of defeat denies the thesis of defeatism put forward at the beginning of the paper. Instead it says that reasoning is non-monotonic, which is precisely what I would like to advance.

[71] After much of the work of this chapter was done I was made aware of Kotzen (2019), which aims to provide a precise, formal account not only of defeat generally, but one which unifies various notions of defeat. I think that Kotzen's account is much more sophisticated than any other such accounts out there. As far as I can tell, the sorts of arguments he advances fall within the scope of my argument. They are mainly probabilistic, hence some of the examples I use could be leveraged against his account. Namely, that he is not able to properly locate what counts as a defeater.

[72] I do not have the space here, but another area where we can see something similar—that is, cases where a set of considerations supports some conclusion, but an expanded set fails to, in virtue not of any particular consideration, but of how those considerations hang together—is via examples from "AGM belief revision" (Alchourrón et al., 1985; Hansson, 2017). AGM belief revision aims to capture the way that we revise our beliefs in light of evidence. We start with a belief base Δ which gives rise (via some closure-like operation, with a number of possible constraints) to a belief-set $\Delta' \supseteq \Delta$. A new piece of information p may be a defeater because $\neg p$ is in Δ'. So we can say, for example that Δ supports q but the addition of p makes it such that $\Delta \cup \{p\}$ no longer does (its closure does not contain q). But this is because the belief *set* associated with Δ contained $\neg p$ and that associated with $\Delta \cup \{p\}$ does not contain $\neg p$. Supposing that $\neg p$ figured in securing q, it's clear how p is a defeater of this. But we can describe cases where there is no such clear-cut relationship between the additional considerations and a proposition which is removed from the belief set in light of revision.

[73] A potential objection is that the defeater is the way that the new information changes the weight of the previous considerations. In particular, the new information weakens the strength of the reason that rain in the mountains and snow in the mountains each give for going to the mountains. But the new information does the exact same thing for the rain and snow considerations in the city. Further, we could easily alter the original utilities to avoid describing things in this way (suppose that rain and snow had a negative outcome, so that the new information actually weakened the reason against).

2.3.3 Beyond Defeat: Substructural Features of Implication

This section has argued that we cannot account for "defeasible reasoning"—the phenomenon in which Γ provides reason for p, but Δ ($\Delta \supsetneq \Gamma$) might not—in terms of the idea of "defeat". This is because often the explanation for why Γ provided reason but Δ does not has nothing to do with any considerations in Δ, but rather with how the considerations hang together; how they are structured. The assumption that any features of Γ's structure will be preserved in Δ (at least with respect to consequence) is called monotonicity. The denial of this is called non-monotonicity. If we cannot capture defeasible reasoning in terms of defeat, however, then it seems inapt to use the label "defeasible". Instead we should understand such reasoning as *non-monotonic* and opt for an understanding that emphasizes this substructural feature of consequence rather than the way in which various considerations might interact.

2.4 Conclusions

In this chapter I put forward a theory of meaning in which the meaning of a sentence is understood as that sentence's contribution to good implication. I did this while also understanding such implication to be possibly radically substructural—i.e. non-monotonic, non-contractive, non-transitive, and/or non-reflexive. Despite the utter lack of constraints on meaning, I provided a formally tractable semantics for sentences understood in this way based upon the phase space semantics for linear logic introduced by Girard (though with some innovations of my own). I also showed that we can define standard propositional logical connectives within the semantics and that the sound and complete proof rules (in the form of a sequent calculus) for these connectives are simple and familiar.

Next, using this tractable proof system, I was able to show that the system can also help us precisify a topic of interest to inferentialists: logical expressivism. Recall that this is the idea that what distinguishes logical connectives is their ability to express features of implication. I showed that not only can we precisify what this means, but that with the help of this proof system, we can show that the logic constructed is expressive in exactly this sense: it can express *exactly* what follows from what. Following this, I also showed another sense in which a logic might be said to be expressive: namely, in terms of its ability to express *structural* features of implication, by which I mean the structural rules we relaxed in the first part of this chapter. To do this, I introduced propositional connectives, which allowed us to express in the object language that an implication obeyed this-or-that structural rule. This is not only an interesting result in the project of precisifying logical expressivism, but also shows that we can introduce connectives to recapture the structural rules that we relaxed. I believe that this is further evidence of the tractability of the system: that we can show ways of recapturing traditional conceptions of content. One such recapture involved introducing a connective for "literally" ($\boxed{L}$), which marked in the object language where an implication obeyed all structural rules considered in this chapter. I proposed that this connective recovers those aspects of a sentence's meaning which were taken to be definitive of its content by traditional semanticists.[74] This contribution—the construction of an operator for "literally"—is not just implicitly interesting, but will be of particular interest in Ch. 4, where literal meaning is discussed.

Finally, I argued that approaching substructural implication in the way I am suggesting may be required for making sense of that substructurality. In this chapter, I suggested that we understand these features of implication as part of the very content of a sentence, but often substructural logics are constructed on top of ordinary semantic notions. I argued that doing this for non-monotonic logic—by understanding the non-monotonicity in terms of defeasible reasoning and thus giving an account of it in terms of defeaters—will not work. If $\Gamma \vdash p$, but $\Delta \nvdash p$ (for some $\Delta \supsetneq \Gamma$), then we cannot rely on the idea that there is a defeater $d \in \Delta$ that explains this phenomenon. My argumentative strategy first focused on the idea of understanding defeat in terms of a relationship between a reason and a particular consideration (the defeater). I argued that there is no unified explanation for why the presence of some further consideration may *defeat* this relationship. Next, I focused on attempts to understand defeat in terms of entire justificatory verdicts. In the weakest account

[74]At least for sentences which have literal meanings. Nothing *requires* that they do and I make no such stipulation that they do.

I considered, we say simply that Δ contains a defeater (without having to specify what that defeater is). But as I argued, there are cases where there is no consideration in Δ that plays the role of a defeater, instead we should understand the change in overall justifcatory verdict in terms of a change to just that: an additional consideration might simply change the way all of our considerations hang together, with the result that a previously justified conclusion is no longer justified. Understanding things in this way, however, is just to acknowledge that implication is non-monotonic (substructural).

In the next chapter, I make more precise the idea of *pushing* the substructurality of implication *all the way down into the content*. That is, I try to make more precise what distinguishes my approach from the sorts of semantics used for other substructural logics. In doing so, I articulate what I call the *assumption of structurality* (to be explained there).

3.0 The Structural Assumption

In the previous chapter I outlined a theory of "substructural content". This is content where its potential to stand in substructural relations of implication—i.e. relations which potentially violate monotonicity, transitivty, contraction, and reflexivity—is built into the content at a "deep" level: I advocated simply understanding the content in terms of its role in good implication and good implications, as I understand it, can be substructural. In addition to putting such a theory of content on the table, I hoped to do two additional things in the previous chapter. First, I hoped to show that despite the lack of constraints on relations of implication, the notion of content I was advancing was completely tractable: we were able to construct a formal semantics, recognizable logics with extremely simple proof systems, and I proved a representation theorem for the logic. Second, I hoped to give some motivation for approaching substructural implication in this way. In particular, I argued that accounting for the non-monotonicity of consequence in terms of something like "defeasible reasoning"—where this is understood in terms of the notion of a defeater—fails to bear fruit. It fails to bear fruit because the target phenomenon (non-monotonic reasoning) often fails to yield a consideration that we can label the defeater (let alone a systematic way of labeling such a consideration).

This chapter seeks to fill in some of the details of my theory of content and to get clearer, for example, on why talk of non-monotonicity is to be preferred to talk of defeasibility. To accomplish this, I use two idioms throughout the chapter, which I hope will shed light on the phenomenon of substructural content. The first idiom involves speaking of **pushing the substructurality all the way down into the content**—or pushing a particular substructural feature, e.g. non-monotonicity all the way down into the content—by which I mean accounting for the substructural nature of implication at the most basic level of a sentence's content. What this idiom is meant to rule out is accounting for the substructural nature of implication by appeal to a "deeper" level of content. For example, in one context a sentence p implies q; but, in another context, the sentence p does not imply q. To get around this, we suppose that p expresses several contents. In the first context, the content associated with p implies q, and importantly *that* content, if it could be isolated, implies q regardless of what considerations are added. The second content does not do so. I claim that the desire to postulate such further contents is motivated by **the assumption of structurality**. The assumption of structurality—the second idiom I intend to deploy—is simply the assumption that if there is a violation of a structural rule of implication, then this must be explained away by appeal to a level of content at which such violations cannot occur

In short: the *assumption of structurality* and the suggestion to *push the substructurality all the way down into the content* are opposed to one another. Pushing the substructurality all the way down into the content involves taking the holistic nature of consequence and embedding it in the very foundation of a sentence's meaning. The assumption of structurality, by contrast, is the assumptions that there must be a basic, logically regimented layer of content that behaves fully structurally (i.e. does not stand in relations of implication that involve violations of structural rules). Do we require a foundation of atomic, logically regimented contents in order to build up "deviant" inferential behavior, or can we simply understand such contents holistically in terms of such behavior?

I try to get clearer on what this opposition entails and thus on the view I am pursuing by examining arguments for (and against) the possibility of various substructural logics. In particular, what I hope to show that is that a large number of approaches to substructural logic (despite their very different motivations and formal settings) fail to give up the assumption of structurality. Because of this, I argue that they do not therefore *take the substructurality of consequence seriously*.[1] In particular, I investigate four different structural rules in this section, MONOTONICITY (MO), CONTRACTION, TRANSITIVITY (or: Cut), and REFLEXIVITY (RE). MO and Contraction have a left- and a right- rule associated with them (since they each only involve sentences on a single side of the turnstile). See Figure 4.[2]

[1]This is not necessarily an indictment of those accounts. Many philosophers and logicians would be happy to hear me report this. Rejecting a structural feature is taken by many to be akin to the suggestion that we revise a fundamental, constitutive rule of thought. If I'm offering an argument that they fail to do that, then I'm confirming what they set out to do.

[2]It's possible to write MO as a single rule since you can treat either Δ or Λ as empty in e.g.:

$$\frac{\Gamma \vdash \Theta}{\Delta, \Gamma \vdash \Theta, \Lambda}$$

Nothing significant hinges on this difference in my presentation.

$$\frac{\Gamma \vdash \Theta}{A, \Gamma \vdash \Theta}\text{ L-MO} \qquad \frac{\Gamma \vdash \Theta}{\Gamma \vdash \Theta, A}\text{ R-MO}$$

$$\frac{A, A, \Gamma \vdash \Theta}{A, \Gamma \vdash \Theta}\text{ L-Contraction} \qquad \frac{\Gamma \vdash \Theta, A, A}{\Gamma \vdash \Theta, A}\text{ R-Contraction}$$

Figure 4: Structural Monotonicity and Contraction

$$\frac{\Gamma \vdash \Theta, A \qquad A, \Gamma \vdash \Theta}{\Gamma \vdash \Theta}\text{ Shared-Cut} \qquad \frac{\Gamma \vdash \Theta, A \qquad A, \Lambda \vdash \Delta}{\Gamma, \Lambda \vdash \Theta, \Delta}\text{ Mixed-Cut}$$

$$\frac{}{\Gamma, A \vdash A, \Theta}\text{ CO} \qquad \frac{}{A \vdash A}\text{ RE}$$

Figure 5: Structural Transitivity and Reflexivity

Transitivity and Cut have a number of different forms.[3] Generally speaking many of these forms are equivalent given MO and Contraction,[4] but since I am investigating these rules as well, we ought not to assume any particular form. I'll outline two forms of each of them for simplicity: RE refers to *simple reflexivity*, i.e. of individual sentences. CO (short for: containment) refers to reflexivity with the addition of arbitrary premises and conclusions. See Figure 5.

For all of these structural features, I try to show that many of the extant accounts out there postulate some further "deeper" level of content which is fully structural. Because of this they don't take the substructurality seriously.

While the main contribution of this chapter is to gain some insight into what substructural content is and what is required to accomodate such content, I'd like to highlight three additional, potentially separable contributions that I make.

1. In the first section, I examining two accounts of non-monotonic reasoning. In both cases,

[3]A great overview regarding transitivity is provided by Ripley (2018). See Restall (2002) for an overview on the same and the other rules I discuss.

[4]I mean that shared and mixed versions of cut are equivalent. There are still differences to be found between Set and Fmla versions of each.

the non-monotonicity of reasoning is accounted for in something besides the content. My suggestion to instead understand that substructural behavior as a feature of the content itself opens up logical space to address a number of issues in the philosophy of language. One such issue—radical contextualism—is discussed in the next chapter.

2. In the second section, I examine various accounts of non-contractive consequence. I spend more time on contraction than I do on other structural features. This is because contraction is often thought to be the most difficult structural feature to imagine giving up (at least of the ones I survey). For this reason, I wanted to provide the most thorough examination of the motivations for abandoning it that I could, and thus to spend more time justifying how I think of non-contractive content. An additional, potentially separable, contribution that my work provides is a better understanding of the relationship between the dominant substructural solutions to paradox, by which I mean non-transitive, non-reflexive, and non-contractive solutions. I argue that in thinking clearly about violations of contraction as I do, we can see a central feature of these violations—what I call multi-facetedness—plays a central explanatory role in paradoxical reasoning.
3. Finally, in the last section, I examine both non-transitive and non-reflexive consequence relations. Part of what my examination helps to bring out is the relationship between these different structural features. In particular, the ways that these four structural rules (or violations thereof) relate to one another.[5] One, potentially separable, insight that is gained from appreciating these relationships is a deeper understanding of the relationship between non-reflexive and non-transitive solutions to paradox.[6]

[5] See the table on p. 112.

[6] I can put the point crudely by saying that they are related to each other in the same way that subvaluational and supervaluational accounts are (where the former corresponds to non-transitive solutions and the latter non-reflexive). I don't claim that non-transitive approaches are subvaluational (nor that non-reflexive are supervaluational)—I think this claim is inapt—which I why you'll also find that I don't make much effort to define either of these terms. References to either appear only in footnotes.

3.1 Monotonicity

This section puts forward an account of what I'll call "non-monotonic content". Put simply this is content (or the possibility of such content) that may figure in non-monotonic inferences. The goal is to investigate what happens if (or whether) content itself is understand to be sensitive to violations of monotonicity (rather than e.g. sentences or expressions of such content).

We might think, for example, that the implication from "s is a bird" to "s can fly" is good, but from "s is a bird" and "s is native to Antarctica" to "s can fly" is bad. A typical account given might make the following move: the word "bird" is context-sensitive. In the first invocation, "bird" refers to a class of animals that can typically fly. But in the second invocation which includes the additional premise that "s is native to Antartica", "bird" refers to something different (namely a class of animals that typically cannot fly). This is to say that while the *sentential* behavior is non-monotonic, the *content* of those expressions is not. For, of anything that is predicated "bird_1" (the first invocation) it will follow of that thing that it can fly; likewise, of anything that is predicated "bird_2" (the second invocation) it will not follow. My suggestion is to ask whether (or to see what happens if) we understand the content, rather than e.g. the sentence which expresses that content, to be non-monotonic in this way.

First, I'll need to provide some preliminaries concerning what I mean by "content". In the previous paragraph, I appealed to at least three distinct levels at which we might understand content.[7] The first (i) was the implicational level, where we can most immediately see the non-monotonic implication relation. What is communicated at this level by "s is a bird" is determined by more than this sentence, as additional sentences may defeat or otherwise interact in complicated ways with e.g. an assertion of that sentence. Beyond this (ii) are the sentences themselves. It is by appeal to this level that the *same thing* can figure in both an implication to "s can fly" and simultaneously figure in a non-implication *to that very same thing*. Finally (iii) there are the actual propositions that these sentences pick out. In

[7] Compare accounts given by (Recanati, 2004; Cappelen and Lepore, 2008). A more nuanced discussion of these issues is presented in the next chapter.

the above "s is a bird" ranged over two distinct propositions: "s is a bird$_1$" which implies that "s can fly" and "s is a bird$_2$" which does not imply this. Which of these propositions the sentence picks out was thought to be determined by what auxiliary commitments or sentences are floating around. The presence of "s is from Antartica" changes *in what sense* s is a bird and (the thought goes) thus changes the content of that sentence.

According to an account like this, the defeasibility of the above inference is apparent at the first (i) level and counts as an instance of *genuine* non-monotonicity because of the second (ii) level (where e.g. "s is a bird" is the same in both implications). But what actually explains what follows from what is explained at the third (iii) level, which functions otherwise in quite traditional ways. In other words, non-monotonicity is sorted out in (i) and (ii), by disambiguating the phenomena. This can be done in a number of different ways (the example above suggests only one such way). What I want to investigate in this section, however, is whether there is a way to understand the non-monotonic behavior of something like "s is a bird" to be a feature of the *content* itself, which in the previous example corresponded to (iii). This requires that the content of "s is a bird" figures univocally both when it implies that "s can fly" and when (together with a further premise) it does not imply this, and the sense in which these are (not) implied (i.e. the sense of "follows from") must similarly be univocal.

Because of the number of moving parts that such accounts deploy it is somewhat difficult to say exactly what this amounts to. Divergent accounts of content, or of inference, or of defeasibility, make it difficult to say exactly how my proposal differs (or even to locate a single feature of such accounts that I wish to distance myself from). Accordingly, I will proceed as follows. I consider two accounts of non-monotonic reasoning and explaining how they fail to push the non-monotonicity *all the way down* into the content.[8] In particular, I look at (i) a preferential models approach, which (as in the above example) does *not* account for non-monotonicity by understanding a sentence like "s is a bird" univocally. I also (ii) look at accounts of default logic, which I claim do not treat "follows from" in a univocal

[8]I do not present an exhaustive characterization of non-monotonic accounts of implication. I do in fact think that most other accounts in the literature have similar features, but I do not argue that here. Missing from this section are so-called "argument based approaches"; see e.g. Pollock (1987, 1991); Batens (2004). For surveys and further citations see Strasser and Antonelli (2018); Koons (2017). I also do not discuss AGM belief revision (Alchourrón et al., 1985) nor probabilistic approaches.

way. I should emphasize now (as I try to throughout) that I do not wish to claim that preferential models or default logic are "wrong". Rather, I wish to claim that they fail to understand the *content itself* as suitable for non-monotonic inference. As accounts of e.g. reasoning about normality or default arguments this is fine, but as accounts of non-monotonic inference generally, I think we can do more justice to the phenomenon.

The main lesson: many accounts of non-monotonic implication start with more-or-less traditional semantic stories concerning the *meaning* of e.g. sentences (or: of e.g. "follows from"). Either something strange is done at the *pre*-semantic level (so that the same sentence may receive divergent contents); or, something strange is done at the *post*-semantic level, by e.g. doing something strange to the definition of model-theoretic entailment. The result of this is that it's misleading to say that the more-or-less traditionally understood semantic interpreteds (e.g. sentences) are interpreted univocally as they figure in substructural relations.

3.1.1 Preferential Models

In an influential paper, Kraus et al. (1990) introduce what they call "preferential models" to capture non-monotonic reasoning. For example, the implication that "if something is a bird, then normally it flies" it meant to be captured by their system. The resulting consequence relations they generate are non-monotonic, but transitive and obey what is called cautious monotonicity.[9]

Definition 3.1.1 (Preferential Model)**.** A *preferential* model W is a triple $\langle S, l, \prec\rangle$, with S a set of "states", l a function from states to worlds (so that each state is assigned a world), $l : S \mapsto \mathscr{U}$. $\prec$ is a strict partial order (irreflexive and transitive) on S that satisfies the further condition that for all $\alpha \in \mathcal{L}$ (sentences in the language) and all states $T_\alpha \subseteq S$ such that $t \in T_\alpha$ iff $l(t) \vDash \alpha$ (i.e. α is true in word $l(t)$), we have that for all $t' \in T_\alpha$ either t' is

[9]Gabbay (1985) has argued that all consequence relations—*a fortiori* all non-monotonic consequence relations—must satisfy transitivity and cautious monotonicity (the condition that $A \mid\!\sim B$ (A defeasibly implies B) and $A \mid\!\sim C$ means $A, B \mid\!\sim C$). Hlobil (2016) offers a clever argument on inferentialist grounds for resisting part of this conclusion. Brandom (2018a) offers further philosophical arguments on inferentialist grounds for resisting the conclusion. The issue is slightly more complicated in my setting because of the use of multi-sets and multiple conclusions. See Arieli and Avron (2000) for a discussion of some of the issues involved with multiple conclusions and cautious monotonicity.

minimal with respect to $\prec$ or there exists $t \in T_\alpha$ with $t \prec t'$.[10]

This might seem complicated, especially concerning "what a sentence means" (or its semantic interpretant). Their picture allows for two levels of semantic meaning. The meaning of the sentence can still be understand as the set of worlds on which it obtains. But there's a super-structure of ordered states on top of that. The following definition will make it clear how particular sentences are understood in particular contexts.

Definition 3.1.2 (Preferential Consequence Relation $\mid\!\sim_W$)**.** We say that $\alpha \mid\!\sim_W \beta$ iff for all minimal $t \in T_\alpha$ (henceforth I'll write T^-_α to refer to the set of minimal states), we have $l(t) \vDash \beta$.

In other words, the minimal α-states are all β-states means that $\alpha \mid\!\sim \beta$. What a sentence means as it appears in a given implicational context, in other words is nothing other than the worlds specified by its $\prec$-minimal states. This means that if $\alpha \mid\!\sim_W \beta$ we should understand α as having the content given by the states T^-_α (and thus by the worlds assigned to those states by l). This is unsurprising since those states are meant to specify "normal conditions" for α. If $\alpha =$ "Tweety is a bird", then T^-_α will consist of states where "Tweety" is flier (hence if $\beta =$ "Tweety can fly", we get $\alpha \mid\!\sim_W \beta$).

Next, suppose we want to inquire what happens when we add γ as an additional premise (suppose $\gamma =$ "Tweety is from Antarctica"). Then we must inquire whether $\alpha \wedge \gamma \mid\!\sim_W \beta$.[11] Per our definitions, we want to know whether for all $t \in T^-_{\alpha\wedge\gamma}$ we have $l(t) \vDash \beta$. Intuitively it should be possible that there is a counter-example (i.e. a t' such that $l(t') \nvDash \beta$). This is possible since, while $T_{\alpha\wedge\gamma} = T_\alpha \cap T_\gamma$, because of how '$\prec$' is defined it does not follows that $T^-_{\alpha\wedge\gamma} = T^-_\alpha \cap T^-_\gamma$ (i.e. it doesn't follow that the minimal sets are also in the intersection). It may be that a is minimal in T_α, c is minimal in T_γ though neither is contained in the other. Further it happens that $d \in T_\alpha \cap T_\gamma$ (though is not minimal in either T_α nor in T_γ). Further it may be that d is minimal in $T_\alpha \cap T_\gamma$. If, further, $l(d) \nvDash \beta$ (which is certainly not ruled out by anything I've stipulated) then we have our result. Figure 6 presents an example of what this might look like. In the figure squares represent states. Lowercase letters represent

[10]This is related to what is standardly called "the limit assumption" in the logic of conditionals (Lewis, 1973; Stalnaker, 1968). See also Kraus et al. (1990, p. 182).

[11]KLM define consequence relations as relations of single sentences to single sentences, hence (non-)monotonicity must be formulated with logically complex sentences in their setting).

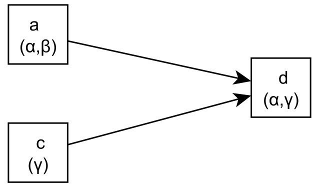

Figure 6: KLM Preferential Model Example

which state it is. Greek letters enclosed in parentheses indicate which of α, β, γ are true in the world assigned to that state by l. The arrows connecting boxes represents $\prec$.

What is significant is that while we considered $\alpha \mid\!\sim_W \beta$, α clearly had the content specified by T^-_α. But what content does α bear in $\alpha \wedge \gamma$? Clearly it must be $T^-_{\alpha\wedge\gamma}$. This is to say that the world specified by $l(a)$ no longer figures as part of the content of α, but $l(d)$ now does (where it had not prior). *Prima facie* this is perhaps how it should be. When we say "Tweety is a bird", a range of "normal situations" comes to the fore. When we add that "Tweety is from Antartica" that range is altered. What we mean by "... is a bird" is suitably changed. On a naive picture of content, there is nothing wrong with this.

What I am trying to advocate for is a picture of content in which "bird" means the same thing in both utterances. We may reason differently, but that's not because we have changed the content of "bird", it's because what "bird" means demands that.

3.1.2 Default Logic

The second sort of example I wish to look at treats the *semantic content* of e.g. bird, flies, penguin, and so forth univocally, but suggests divergent treatment of "follows from". To do this I will examine a particular account of default reasoning found in Horty (2012).[12]

First I define what is called a "fixed-priority default theory".[13]

[12] Default *logic*—as an account of default reasoning—was first put forward as a kind of fixed-point approach to non-monotonic logic by Reiter (1980) (see also Antonelli (1999); Makinson (1994)). Reiter's account contains enough sophistication for my purposes here, but I instead choose to follow Horty because of its relevance and because its increased sophistication provides some nice parallels to the preferential models example in the previous subsection. For more on Horty's account see also Horty (2001, 2007a,b).

[13] Horty's account relies on two innovations over earlier accounts: its use of a priority relation between

Definition 3.1.3 (Fixed Priority Default Theory)**.** Let $\mathcal{W}$ be a set of propositions, $\mathcal{D}$ be a set of default rules and $<$ a strict partial ordering (i.e. a transitive irreflexive ordering) on $\mathcal{D}$. Then call $\Delta = \langle \mathcal{W}, \mathcal{D}, < \rangle$ a fixed priority default theory.

This definition requires a comment. A default rule is a kind of inference ticket. Horty introduces defaults in a non-orthodox manner (though I think his method is more perspicuous for my purposes). According to more traditional default logicians, a default rule consists of (roughly) three parts, some known sentence(s) σ, a conclusion γ and some kind of constraint (or constraints) β that determine whether this default rule is applicable or not:

$$\frac{\sigma : \beta}{\tau}$$

Default logic proceeds by investigating various *extensions* of a default theory. Where an extension is the logical closure of a default theory together with whatever can be concluded from the default rules. Reiter looked, for example, at the intersection of extensions. Horty takes a slightly different approach, but the idea is similar: there is some proper set of default rules that (depending on the composition of $\mathcal{W}$) an ideal reasoner would make use of.[14]

For now, we shall think of default rules as connecting two propositions $X \mapsto Y$. So if the previous is the default rule δ then we can write that $X = \text{Prem}(\delta)$ and $Y = \text{Concl}(\delta)$. Next, we next define a scenario $\mathcal{S}$ as a subset the default rules $\mathcal{D}$. As shorthand we can write $\text{Concl}(\mathcal{S})$ for $\{\text{Concl}(\delta) | \delta \in \mathcal{S}\}$ (and similarly for $\text{Prem}(\cdot)$). Likewise, we can understand (intuitively) an extension $\mathcal{E}_{\mathcal{S}}$ of a default theory Δ *generated by that scenario* as the logical closure of $\mathcal{W}$ together with the conclusions generated from $\mathcal{S}$:

$$\mathcal{E} = Th(W \cup \text{Concl}(\mathcal{S})).$$

Here $Th(\cdot)$ means the logical closure. I will also use $\vdash$ in this subsection to indicate logical consequence. While it is not required, default logicians often mean here classical consequence and classical closure (nothing in any of the examples here hinges on this).

What we want next then is what makes a scenario a good one (so that we can know which extensions to pay attention to: presumably those generated from good scenarios).

---

default rules and its use of a non-fixed relation. For the purposes of this section, I only introduce the first of these two things. See Horty (2012, pp. 22ff.).

[14] To perhaps oversimplify: Horty builds the constraints into scenarios rather than the default rules.

Intuitively the stable scenarios are those that include exactly those default rules that are actually triggered by the propositions in $\mathcal{W}$ and none which conflict with or defeat anything in the scenario.

Definition 3.1.4 (Stable Scenario)**.** Given a default theory $\Delta = \langle \mathcal{W}, \mathcal{D}, < \rangle$ and scenario $\mathcal{S}$, we say that $\mathcal{S}$ is proper iff[15]

$$
\begin{aligned}
\mathcal{S} = &\{\delta \in \mathcal{D} : \mathcal{W} \cup \mathrm{Concl}(\mathcal{S}) \vdash \mathrm{Prem}(\delta)\} && \text{(triggered)} \\
&\setminus \big(\{\delta \in \mathcal{D} : \mathcal{W} \cup \mathrm{Concl}(\mathcal{S}) \vdash \neg\, \mathrm{Concl}(\delta)\} && \text{(conflicted)} \\
&\cup \{\delta \in \mathcal{D} : \exists \delta' > \delta((\mathcal{W} \cup \mathrm{Concl}(\mathcal{S}) \vdash \mathrm{Prem}(\delta')) \text{ and } (\mathcal{W} \cup \{\mathrm{Concl}(\delta')\} \vdash \neg\, \mathrm{Concl}(\delta))).\}\big) \\
& && \text{(defeated)}
\end{aligned}
$$

Following Horty I have labeled each "clause" in the definition as "triggered" (for those rules which are *triggered* by the scenario given the theory), "conflicted" (for those rules which conflict with those in the scenario given the theory), and "defeated" (for those rules which are defeated in the scenario given the theory).

Intuitively a proper scenario (relative to a default theory) is one that contains exactly the default rules which it (together with the default theory) triggers and contains *none* of the rules which either conflict with it (in the sense that the scenario including its defaults logically imply the negation of the conclusion of another default) or which the scenario defeats (i.e. a less preferable$_<$ rule which conflicts with a rule that is triggered by the scenario).

Now, we are in a position to define extension.

Definition 3.1.5 (Extensions)**.** Given a default theory $\Delta = \langle \mathcal{W}, \mathcal{D}, < \rangle$ and *proper* scenario $\mathcal{S}$. An extension of Δ (based on $\mathcal{S}$) is defined as:

$$\mathcal{E} = Th(\mathcal{W} \cup \mathrm{Concl}(\mathcal{S})).$$

Whether a theory has multiple (or no) extensions will be a function of how many proper scenarios it has (if it has any). This complication is important for getting a proper definition of consequence. Different default logicians might advocate for different approaches here. For example, we might have it that a sentence is a consequence of some initial set of propositions

[15] I skip over many of the details of Horty's account since the subtleties are unimportant for my arguments here. They concern what are, by Horty's lights: "aberrant default theories" (Horty, 2012, p. 32).

$\mathcal{W}$ if it figures in *some extension* (the credulous approach). Or we might have it that we only want to call consequences those sentences that are in *all extensions* of a set of sentences $\mathcal{W}$ (the skeptical approach). Interesting though this subtletly is, I will only consider examples (below) with unique proper scenarios, such that both approaches will yield the same approach. Thus, relative to a set of defaults $\mathcal{D}$ we define a non-monotonic consequence relation. My definition here diverges from Horty and the tradition by not taking any stance on the previous paragraph.

Definition 3.1.6 (Default Consequence (Super-Careful))**.** Let $\mathcal{D}$ be a set of defaults ordered by the strict partial order $<$ and let $\mathcal{W}$ be a set of sentences. Then we write (pronounced $\mathcal{W}$ default-implies A):

$$\mathcal{W} \mid\!\sim_{\mathcal{D}} A,$$

iff the default theory $\Delta = \langle \mathcal{W}, \mathcal{D}, < \rangle$ contains a *unique* proper scenario $\mathcal{S}$ with $A \in Th(\mathcal{W} \cup \text{Concl}(\mathcal{S}))$ (i.e. the unique extension of $\mathcal{W}$).

To see how this will generate a non-monotonic consequence relation. Suppose that Δ contains the following *two* rules: $B \mapsto F$ and $P \mapsto \neg F$ (δ_1 and δ_2 respectively). (where $B =$ "Tweety is a bird", $P =$ "Tweety is a penguin", and $F =$ "Tweety is a flier"). Since Tweety's being a penguin should overrule default rules where he is only a bird (the thought goes), $<$ will be exhausted by one relation: $\delta_1 < \delta_2$. Then, we want to establish that $B \mid\!\sim_{\Delta} F$, but $B, P \mid\!\sim_{\Delta} \neg F$.

To do this, we need to set-up two default theories (since we have two separate set of propositions to consider $\{B\}$ and $\{B, P\}$):

$$\Delta_1 = \langle \{B\}, \mathcal{D}, < \rangle \quad \text{and}$$
$$\Delta_2 = \langle \{B, P\}, \mathcal{D}, < \rangle.$$

Next, there are only four possible scenarios we must consider (all members of $\mathcal{P}(\mathcal{D})$). Because δ_1 and δ_2 conflict (their conclusions are negations of one another) we can immediately rule out $\mathcal{D}$ as itself a proper scenario. Further, since the sets of propositions are guaranteed to trigger at least one default for each theory, we can likewise rule out $\emptyset$ as a possible proper scenario.

As it turns out, $\{\delta_1\}$ is the unique proper scenario for Δ_1. It is proper since its premise is triggered. It is unique because $\{\delta_2\}$ is not proper for Δ_1 because it doesn't contain all triggered rules. Similar considerations show us that $\{\delta_2\}$ is the unique proper scenario for Δ_2.

It should therefore be obvious that since

$$F \in Th(\{B\} \cup \mathrm{Concl}(\delta_1)), \quad \text{and}$$
$$\neg F \in Th(\{B, P\} \cup \mathrm{Concl}(\delta_2))$$

that:

$$B \mid\!\sim_{\mathcal{D}} F, \quad \text{and}$$
$$B, P \mid\!\sim_{\mathcal{D}} \neg F.$$

In particular, we have that $B, P \not\mid\!\sim_{\Delta} F$.

In contrast to the previous example the "content" of B, P, F was univocal accross these inferences. That is, the sense in which Tweety is a bird or (not) a flier does not change. But what *does* change is the sense of "follows from" in ascertaining whether F follows from B or from B, P. While we said that they followed from the set of defaults $\mathcal{D}$, in fact it was different (sets of) rules that did the work. *Prima facie* this is perhaps how it should be. When we say that "Tweety is a bird" some default rules come to the fore and when we add that "Tweety is a penguin" that range of default rules is suitably changed. This is to say that when we determine *whether it follows* that "Tweety can fly" we are consulting *different rules* and the sense in which it *follows* that Tweety can (not) fly is thus suitably altered (because the manner of determining whether that is so is suitably altered).

I do not mean to deny that there is a perhaps perfectly innocent sense in which we can say "follows from" across both cases (in the same way that there might've been a way to say "... is a bird" in the previous section). But I don't think this is the most natural way of understanding the situation. I also do not wish to claim that there is something wrong with this account. As an account of default reasoning, it seems to get things right.

Rather, what I am trying to advocate for is a picture of content on which we get the relevant non-monotonic behavior not because we have shifted the behavior of "follows from" but because the content of "Tweety is a bird" by itself behaves in this way.

In this section, I gave two examples of non-monotonic logics that do not treat non-monotonicity as a feature of the content of a sentence; that is, they do not push the non-monotonicity *all the way down* into the content. In particular, preferential-models account for non-monotonicity by implicitly taking those sentences to have different contents (dependent upon what further assumptions are active). I also contrasted my account with default logic which—though it understands the content which figures in defeasible implications univocally—implicitly invokes different notions of follows from to account for that non-monotonicity (dependent upon what further assumptions are active). As I've tried to emphasize: I don't think there is anything necessarily wrong with such accounts, but I also don't think they do as much justice to the general phenomenon of non-monotonic reasoning as they can. I hope to have cleared space for thinking about what sort of *content* can figure in non-monotonic reasoning.

My own account treats the non-monotonicity of reasoning as part of the content of sentences which stand in those relations. Thus, in contrast to the two approaches surveyed here, we can say the non-monotonicity is pushed all the way down into the content.

3.2 Contraction

Next, I would like to discuss contraction. Because contraction is the hardest structural rule to deny, I try to spend the most time getting clear on why others have denied it and what my own denial entails. The goal, therefore, is to convince the reader that the goodness of the implication:

$$\text{from } p, \Gamma \text{ to } q,$$

is *not always* secured by the goodness of the implication:

$$\text{from } p, p, \Gamma \text{ to } q.$$

If we use '$\vdash$' to symbolize that an implication holds, then I am denying the following:

$$\text{from } p, p, \Gamma \vdash q, \text{ we } \textit{always} \text{ have } p, \Gamma \vdash q.$$

This principle, which allows one to ignore how many tokens of p occur as a premise is called (structural) contraction. I think structural contraction doesn't hold. I also think that the arguments in the literature against it are less than fully convincing. I'm hoping to change that. The main argumentative joint of this section consists in setting up an opposition—between reasons why contraction is obvious and arguments against the obvious—and then showing that the opposition rests on a faulty presupposition. Once we deny the presupposition, there are convincing paths to giving up contraction that don't involve rejecting the obvious.

It is often taken that the following obvious constraint on meaning entails contraction (i.e. that the following argument is valid):

P1: (NO EQUIVOCATION) In a proper accounting of reasoning, sentences must be fully disambiguated. Two occurrences of the same sentence must express the same content in the implications in which they are involved on pain of equivocation.

C: (CONTRACTION) In a proper accounting of reasoning, implications obey contraction: if an implication with two occurrence of A as a premise is good, then the same implication with only one occurrence is good as well.

If this is true, then the only way to deny contraction is to deny **P1:** NO EQUIVOCATION. Steps can be taken to soften this denial, but we must nevertheless allow for a different picture of content.

I argue instead that there is a gap in this argument that once filled allows us to deny contraction without denying the above constraint on content. In other words, I argue that the argument as it stands is invalid and requires a further premise:

P2: (SAMENESS IS REDUNDANT) In a proper accounting of reasoning, if two occurrences of the same sentence each make the same the same contribution to consequence, then the presence of either is inferentially redundant (its removal cannot infirm a good implication).

I argue that if we adopt an inferentialist understanding of meaning, according to which the meaning of a sentence is to be understood in terms of its contribution to good implication, then logical space opens up for us to deny **P2** instead of **P1**. Thus we have a way of denying contraction without denying the obvious constraint expressed by **P1**.

I begin (§3.2.1) by explaining where rejections of contraction originated from as well as the typical reasons given to reject those arguments. The upshot of this section is a number of insights into what sorts of behavior might motivate a rejection of contraction, as well as arguments to keep it. All of the arguments seem to hinge on denial of **P1:** No Equivocation. I suggest that what these non-contractive accounts perhaps are really gesturing at is another feature content may have, namely that content can be "multi-faceted". Endorsing the view that content is multi-faceted need not involve a rejection of **P1**. I close this section by arguing that we shouldn't foreclose the compossibility of denying **C:** Contraction while affirming **P1:** No Equivocation if we allow for content to be multi-faceted. For this to bear fruit, however, I'll have to say more about how this works.

In (§3.2.2), I explain that we can gain such an understanding of multi-facetedness by adopting an inferentialist understanding of consequence according to which the meaning of a sentence is to be understood in terms of its contribution to good implication. On this understanding it is perfectly intelligible that sentences may be fully disambiguated (i.e. satisfying No Equivocation) and so make the exact same contribution in all of the inferential contexts in which they appear. Nevertheless, two instances of the same sentence are *not* inferentially redundant, which is to say that violations of contraction are possible. The formal, inferentialist, semantics developed in the previous chapter should make the philosophical account just rehearsed formally tractable and more precise. I also provide in this context a formal precisification of multi-facetedness and in particular of non-contractive content within this framework (§3.2.3). In summary: a sentence is non-contractive (i.e is involved in a violation of contraction) iff it contains a facet with which it is inferentially productive. This account is applied to the semantic paradoxes (I focus on the liar and Curry) in order to show how this manifests in cases that have long concerned logicians interested in contraction. I close this section by noting a number of explanatory virtues that multi-facetedness has, especially when the account is situated relative to other substructural solutions to paradox.

The upshot is a general account of content which is non-contractive and recognizable as the sort of thing involved in an account of *reasoning* that does not deny the obvious constraint on content in No Equivocation. The account is useful for making sense of not only material violations of contraction (i.e. of sentences which are non-contractive in virtue of material facts concerning their use), but also for making sense of how e.g. a naive theory of truth results in non-contractive content.

3.2.1 What is contraction, where did it come from, and what is contraction?

The goal of this section is to justify the idea that the goodness of the following argument is taken for granted:

P1: (No Equivocation) In a proper accounting of reasoning, sentences must be fully disambiguated. Two occurrences of the same sentence must express the same content in the implications in which they are involved on pain of equivocation.

C: (Contraction) In a proper accounting of reasoning, implications obey contraction: if an implication with two occurrence of A as a premise is good, then the same implication with only one occurrence is good as well.

I do this by explaining how arguments *for* contraction invoke No Equivocation while arguments *against* contraction (i.e. arguments for non-contractive logics) deny it.

Rejection of contraction originates in the literature on paradoxes of self-reference. Curry's Paradox arises for any theory sufficiently expressive to entail that a sentence κ is equivalent (inter-substitutable) with a sentence that says it implies absurdity (or some falsehood):[16]

$$\kappa_q =_{df.} Tr(\kappa_q) \to q,$$

In prose:

If this sentence (i.e. κ_q) is true, then q.

[16]This can happen in a number of ways. For example, the semantic versions of this make use of a truth predicate ($\kappa =_{df.} Tr(\kappa) \to \perp$); set-theoretic versions make use of set-membership: $\kappa =_{df.} \{x : x \in x \to \perp\}$, then the Curry-sentence is: $\kappa \in \kappa$. Curry (1942) proposed the latter. The former were put forward by Geach (1955); Löb (1955). Sometimes it's necessary to set-up a Curry-schema such that each Curry sentence κ_q is defined to imply q (and thus triviality is yielded by the schema rather than $\perp$ and explosion). This is particularly important to notice because while the Liar and $\perp$-Curry seem to involve sentences that yield absurdity, the Curry *schema* only has it that each sentence yields some falsehood. For the most part I limit discussion of Curry to the "particular falsehood" version.

$$
\begin{prooftree}
\AxiomC{$[\kappa_q]$}
\RightLabel{Def.}
\UnaryInfC{$\kappa_q \to q$}
\AxiomC{κ_q}
\RightLabel{MP}
\BinaryInfC{q}
\RightLabel{Cond.-Intro.}
\UnaryInfC{$\kappa_q \to q$}
\AxiomC{$[\kappa_q]$}
\RightLabel{Def.}
\UnaryInfC{$\kappa_q \to q$}
\AxiomC{κ_q}
\RightLabel{MP}
\BinaryInfC{q}
\RightLabel{Cond.-Intro.}
\UnaryInfC{$\kappa_q \to q$}
\RightLabel{Def.}
\UnaryInfC{κ_q}
\RightLabel{MP}
\BinaryInfC{q}
\end{prooftree}
$$

Figure 7: Natural Deduction Presentation of Curry's Paradox

The paradoxical reasoning proceeds as follows. In general if we can show that ψ on the supposition that ϕ, then we will have shown $\phi \to \psi$. So: suppose that κ_q is true. Clearly on this supposition we have κ_q, and clearly we have $Tr(\kappa_q) \to q$ (this is equivalent to κ_q), so by modus ponens we have q. But then we have shown q on the supposition that κ_q, so we have shown $Tr(\kappa_q) \to q$, that is, we have shown that κ_q is true. But if we have both of these, then, once again via modus ponens we have q. Thus we have shown from only the definition of κ_q and some basic facts about the conditional that q. Since q was arbitrary we could define a Curry sentence for every sentence of the language. But then any language which permits Curry sentences is trivial (all sentences are true).

The reasoning here might seem slippery. It's easier to see what's going on, and why we must appeal to contraction if we write it out in a Gentzen style natural deduction (Figure 7).[17] Notice that, strictly speaking, two instances of κ_q are not discharged. This corresponds to our multiple uses of κ_q in the informal deduction. If we allow double discharge (i.e. to discharge both instances of κ_q when introducing the conditional), then we may count the above as a derivation of q. If we do not allow double discharge of assumptions, then the above is at best a deduction of q on the assumption that κ_q. This would be symbolized as: $\kappa_q, \kappa_q \vdash q$, which, of course, does not threaten to trivialize the consequence relation. The connection to the principle of contraction mentioned at the outset can be seen from its sequent derivation (Figure 8).[18]

Curry sentences generated a lot of interest because (i) like the liar they generate problems

[17]In this first section I often suppress $Tr(\cdot)$ for expediency of presentation.

[18]There are a number of ways to avoid the "converse conditional principle" (CCP), but they are all slightly more complicated and amount to the same.

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\kappa_q \vdash \kappa_q}{\kappa_q \vdash \kappa_q \to q}\,\text{Def.}}{\kappa_q, \kappa_q \vdash q}\,\text{CCP}}{\kappa_q \vdash q}\,\text{Contraction}}{\vdash \kappa_q \to q}\,\text{Cond.-Intro.}}{\vdash \kappa_q}\,\text{Def.} \qquad \dfrac{\dfrac{\dfrac{\kappa_q \vdash \kappa_q}{\kappa_q \vdash \kappa_q \to q}\,\text{Def.}}{\kappa_q, \kappa_q \vdash q}\,\text{CCP}}{\kappa_q \vdash q}\,\text{Contraction}}{\vdash q}\,\text{Cut}$$

Figure 8: Sequent Calculus Presentation of Curry's Paradox

for languages which are sufficiently expressive for genuine self-reference,[19,20] but unlike the liar paradox, they don't seem to rely on negation. In fact simply messing with the conditional (as analogously some solutions to the liar mess with negation) seems insufficient. For we can write a generalized version of the Curry paradox.[21]

For example, all that we need to get paradoxical reasoning is a connective '$\odot$' and a sufficiently expressive language to yield a sentence κ_q that allows for the two principles in Figure 9.[22] If '$\odot$' is a a kind of conditional ($\odot\kappa_q = \kappa_q \to q$) then the above is simply Curry's paradox and T is *modus ponens* and C is contraction. But $\odot$ could also be a kind of negation (as in the liar) ($\odot\kappa_q = \neg\kappa_q$) in which case T looks a lot like explosion (*ex falso quodlibet*) and C is contraction (on the right).[23] But $\odot$ could also be a kind of validity predicate or a

[19]Non-contractive solutions to Curry typically reject hierarchichal solutions (Tarski, 1936) to self-reference, type-theoretic (*a la* Russell), or other solutions which seek to limit a naive (i.e. unrestricted) understanding of self-reference (as in e.g. Kripke's 1975 fixed-point solution, or Gupta and Belnap's truth-revisionary solution (Gupta, 1982; Gupta and Belnap, 1993); or the so-called contextualist solutions (Burge, 1979; Barwise and Etchemendy, 1987; Simmons, 1993; Glanzberg, 2001, 2004). This is sometimes put in terms of whether one accepts *naive comprehension* (in the set-theoretic case) or a *naive theory of truth* (in the semantic case). Non-contractive solutions arise in contexts where we wish to allow unrestricted *self-reference*.

[20]Prior (1955) writes:

> "Curry's paradox does not involve the negation but that even Russell's paradox presupposes only those properties of negation which it shares with implication".

[21]Prior (1955); Priest (1994); Shapiro and Beall (2018). Restall (1993) also observes something similar.

[22]Presentation of generalized Curry is lifted from (Shapiro and Beall, 2018). See also Bimbo (2006); Humberstone (2006).

[23]In this case denying contraction on the right looks a lot like a generalization of para-complete solutions to the liar.

$$\frac{\vdash \kappa_q \qquad \vdash \odot\kappa_q}{\vdash q}\,\mathrm{T} \qquad\qquad \frac{\kappa_q \vdash \odot\kappa_q}{\vdash \odot\kappa_q}\,\mathrm{C}$$

Figure 9: Principles for Generalized Curry/Liar Paradox

number of other devices that have emerged in the literature.[24] In all of these cases, however, C shows up as a kind of contraction, so clearly contraction plays a deep role in paradoxical reasoning and thus exploring the possibility of its denial ought to shed some light on what is going on.[25]

3.2.1.1 In Defense of Contraction

Nevertheless, just because it is crucial in constructing paradoxes doesn't mean it ought to be questioned.[26] Diagonalization—the method by which paradoxical sentences are typically constructed—also seems to play an indispensable role, but that doesn't mean we should give it up. Part of what substructural solutions to paradox aim to do is to preserve a *naive* theory of e.g. truth/sets/validity and revise something about our reasoning in paradox. But if the cost of preserving e.g. a naive theory of validity is to radically change what we mean with the turnstile, then it's hard to not see this as simply shooting ourselves in the foot: the

[24] Whittle (2004); Bimbo (2006); Shapiro (2011, 2013); Beall and Murzi (2013).

[25] In the literature on sub-structural solutions to the paradoxes it should be noted that there are three structural rules, strictly speaking, needed to generate paradoxical reasoning. Transitivity (or cut) corresponds to T; contraction to C; in addition reflexivity (or idempotence) is typically used to construct $\kappa_q \vdash \odot\kappa_q$ (French, 2016; Meadows, 2014). There are of course even finer way of dividing things. For example, T may be split into a few different principles some of which seem to involve contraction.

There is a joint in the literature I am also running roughshod over. The focus on structural contraction (as opposed to features of e.g. the conditional which are entailed by structural contraction) gained particular significance in light of the validity-Curry paradox—a sentence which is defined to be equivalent to a sentence which says that the implication from *it* to absurdity (or some falsehood) is valid:

$$\pi =_{df.} Val(\ulcorner\pi\urcorner, \ulcorner\bot\urcorner).$$

See Beall and Murzi (2013). I am generally taking it for granted that giving up contraction offers a unified solution to a number of paradoxes and questioning independent justifications given.

[26] Beall and Murzi (2013) offer giving up contraction as a means to solving Curry, but also note that if this is the *only* reason that can be mustered, then that is hardly a solution.

whole point of the naive theory was that the notion in question requires, on its own, no deep theorizing.[27]

Beall and Murzi (2013) suggest that there's an intuitive problem with dropping structural contraction:

> "If we have assumed a, we are, it would seem, reasoning about a situation in which a is true. Call this situation s. Then , one might argue, surely it should not matter how many times a is used while we reason about s given that a is true in s. On this way of thinking, Structural Contraction would seem to be *essentially* built into ordinary reasoning" (Beall and Murzi, 2013, p. 163).

They suggest that a remedy might be to move away from "certain—standard—conceptions of what validity is". The suggestion is that giving up contraction forces one to supplant ordinary notions of validity and reasoning with something else.[28]

Scharp (2013, pp. 80f.) writes:

> "Imagine someone who says "I accept that the standard model of particle physics, the claim that leptons do not travel faster than the speed of light, and the claim that leptons do not travel faster than the speed of light together entail that neutrinos do not travel faster than the speed of light; however, I do not accept that the standard model of particle physics and the claim that leptons do not travel faster than the speed of light entail that neutrinos do not travel faster than the speed of light." It is hard to know how to even interpret such a claim. In fact, I will bet that many readers had to read back over it several times to make sure it says what it seems to say.[29] The very idea that one token of a sentence has certain entailments, but two tokens of the same sentence have different entailments is profoundly antithetical to the way we think and reason."

Which is precisely to affirm **P2** (that Sameness is Redundant) and thus to affirm the move form **P1** to **C**. Ripley (2015) likewise argues that violations of contraction are incompatible with how we ordinarily conceive of reasoning.

So it seems clear that those places where arguments are mounted in favor of contraction, Sameness is Redundant is taken for granted and it is assumed that the argument from

[27]On the other hand, naive theorists respond by saying that all logicians start with an intuitive notion of "true". So, given the centrality of truth to model-theoretic semantics, revising what "true" means post-hoc is just as problematic.

[28]Beall (2007); Field (2008) suggest moving away from "ordinary" understandings of validity for this reason.

[29]It's of the form:

> "I accept that p, q, and q together entail that r; however, I do not accept that p and q entail r."

$$
\dfrac{
\dfrac{\dfrac{\dfrac{[\kappa_q^1]}{\kappa_q \to q}\,\text{Def.} \qquad \kappa_q^2}{q}\,\text{MP}}{\kappa_q \to q}\,\text{Cond.-Intro.}
\qquad
\dfrac{\dfrac{\dfrac{\dfrac{[\kappa_q^1]}{\kappa_q \to q}\,\text{Def.} \qquad \kappa_q^2}{q}\,\text{MP}}{\kappa_q \to q}\,\text{Cond.-Intro.}}{\kappa_q}\,\text{Def.}
}{q}\,\text{MP}
$$

Figure 10: Natural Deduction Presentation of Curry's Paradox (with Labels)

P1: No Equivocation to **C:** Contraction goes through. To give up contraction requires changing the topic (to something besides reasoning) or fundamentally altering how we conceive of content (rejecting No Equivocation). I'll now turn to accounts that actually defend giving up contraction to see how they do so. As I suggested at the start of this section, my suggestion is that the commitment to the goodness of the argument from **P1** to **C** leads them to reject **P1** (or change the topic).

3.2.1.2 Giving Up Contraction

In light of this, a denial of contraction can't simply be justified in terms of its ability to avoid paradox. So a more general argument is needed. What I argue in the remainder of this section is that non-contractive accounts appeal to the same sort of phenomenon: a certain kind of multi-facetedness that sentences exhibit. They appeal to this by denying **P1** though as I go on to argue this denial is not necessary.

We can see what I have in mind by re-examining the natural deduction proof I offered above for Curry and labeling the distinctive uses of κ_q (Figure 10). There are two distinct ways that κ_q is used in this derivation.[30] In the right assumption of each branch (labeled κ_q^2) it is treated as a sentence variable like any other. It can be used in conditional elimination (MP) if an identical sentence figures as the antecedent of a conditional. The assumption on

[30]NB: $\kappa_q^1 = \kappa_q^2$. The superscripts should be understood as labels for us, and in particular, labels for a position in the natural deduction. They should *not* as distinguishing two sentences. Later in this paragraph I use κ_q to refer to the exact same sentence (not a different sentence with the same content). Formatting restrictions forbids the use of color in dissertations, which I why I use this unfortunate labeling convention here.

the left of each branch (κ_q^1), by contrast, plays a much more special role. This is partially obscured by the level of abstraction featured in the proof tree, but the step labeled "Def." requires us to reason about very special facts concerning the internal structure of κ_q. Namely that κ_q is a special sort of sentence that is equivalent to a sentence which says that it (i.e. κ_q) only if q (for some q).[31,32]

Weber (2014), for example, argues that contraction fails because "a single sentence can be non-self-identical" insofar as occurrences of it may represent distinct aspects of the same sentence (p. 111). In particular, he constructs a hierarchy of validity predicates, V_i such that $V_i(\langle \ulcorner \Gamma \urcorner, \ulcorner \Delta \urcorner \rangle)$ iff $\Gamma \vdash \Delta$ and any V_j that appears in $\Gamma \cup \Delta$ is such that $j < i$.[33] If we can show that this hierarchy reaches a fixed point (which Weber shows) then we can establish the existence of a predicate $\mathbf{V}$ that contains all the validity predicates:

$$\mathbf{V} = \{\langle \ulcorner \Gamma \urcorner, \ulcorner \Delta \urcorner \rangle | \exists i (V_i(\langle \ulcorner \Gamma \urcorner, \ulcorner \Delta \urcorner \rangle))\}.$$

Importantly, Weber denies contraction because two occurrences of the same sentence might be ambiguous. They are "equivalent" in an extensional sense: if one holds, the other holds; but they are grounded in different levels of the validity hierarchy and thus intensionally distinct. Thus, we can avoid contraction on Weber's view precisely by denying something

[31] In what follows I stick as exclusively as possible to substructural logics at the sequent level. Thus, I ignore Hilbert style non-contractive logics (which deny $(A \to (A \to B)) \to (A \to B)$) (e.g. a number of relevance logics). For proposals of this form, see e.g. Restall (1994); Brady (2006); Priest (2006); Field (2008); Beall (2009). Brady's suggestion (as is the suggestion of many relevance theorists who deny this form of contraction) is that just because A secures B does not mean that it secures this sort of connection between A and B. That is: A is playing two sorts of roles in the antecedent: as the sort of thing that can secure the truth of a conditional sentence and as the sort of thing that can secure a particular consequent of a conditional (i.e. B). This itself is a kind of multi-facetedness, which we ought not conceal or equivocate concerning. For example just because A secures $A \to B$ doesn't mean that it is actually sufficient on its own to secure B (that it is relevant in the right sort of way). For example, that it rains secures that the ground is wet ($R \to W$), but *that rain secures wet ground*, is not secured by its raining, it is secured by the laws of physics or perhaps something analytic about rain and wetness. In other words: it is secured by a different *kind* of fact. A sentence must be capable of divergent behaviors to not only secure something further, but also to secure its connection to that further thing.

[32] Even though Shapiro (2015) offers an account against the structural rule of contraction, I feel that it is most easily grouped with the accounts of the conditional in the previous footnote. In short, Shapiro understands $A \vdash B$ as a statement in our metalanguage, namely that $A \Rightarrow B$ is true. $A, A \vdash B$ says $(A \vdash A) \Rightarrow B$ is true. Contraction says $A \Rightarrow (A \Rightarrow B)$ is true only if $A \Rightarrow B$ is true. At any rate, clearly A is playing a different role when it flanks a turnstile (it is a name of a sentence) than it does when it flanks an object language "consequence connective" (it is a sentence). This is, of course, the same sort of multi-facetedness that I gestured to in the previous footnote. Shapiro has legitimate concerns surrounding the validity predicate for doing this that I don't discuss in this paper.

[33] $\ulcorner \Gamma \urcorner$ is the name of the set Γ. This can be set-up in various ways.

like **P1:** No Equivocation. By claiming that two occurrences of the same sentence are not *intersubstitutible*, what we are in effect claiming is that two occurrences of the same sentence need not make the same contribution to the goodness of the implications in which they occur.

Something similar holds for another account of non-contractive consequence found in Zardini (2011, 2019). Zardini cashes out failures of contraction as kinds of "instability" of states of affairs. The liar, for example is unstable because a sentence it "causes" (namely the state of affair of it itself being false) is incompatible with the original state of affairs. Zardini has also noticed a similar kind of multi-facetedness that sentences may exhibit insofar as he acknowledges that a state of affairs may behave inferentially distinctly as a cause and an effect. Ambiguity and equivocation in language are analogous to instability of states of affairs in metaphysics. Zardini, thus, can be seen as making the same kind of move described above: that in order to deny contraction, we have to give up No Equivocation. In the metaphysical case, this means that a single sentence—meant to correspond to a metaphysical state of affairs—might be indeterminate between two separate metaphysical possibilities:[34]

> "The two possibilities can remain open and reality itself unresolved. In other words, it is not that the relevant portion of reality has decided between possibility x characterised by s and possibility y characterised by t while one is still undecided about which of those two SAs characterises reality; rather, one has decided that s characterises both x and y, but the relevant portion of reality is still undecided between x and y. Instability fractures reality into unresolvable possibilities" (Zardini, 2019, pp.180–1).

Essentially, Zardini is invoking a kind of ambiguity for metaphysical states of affairs—such sentences are "double faced" (p. 156). A concrete example might help (borrowed and augmented from decision theory): suppose that you have a rare disorder which causes the cells of your lungs to become cancerous. With no intervention, the chance of cancer is quite high (let's say at least 75%). While there is no known cure, chemicals released from tobacco smoke work to suppress a gene instrumental in causing the cancer to form. Unfortunately,

[34]Zardini's account has faced some technical (Da Ré and Rosenblatt, 2018) and philosophical (Standefer, 2016) objections; in particular, while it's clear that the liar is unstable with respect to itself, Curry requires instability with respect to sentence *type* (since it is only the *type* κ_q that produces triviality—if we allowed a single Curry sentence in, we'd just get a single falsehood). So κ_q isn't unstable with respect to its token effects, but the schema is unstable with respect to all of its effects (or its schematized effect). At any rate, it seems a distinct advantage to be able to say what is wrong with a distinct κ_q without appealing to facts about the schema. Isn't it enough if q is something known to be false?

the risk of developing cancer from smoking itself is unmitigated by the disease. The result is an overall lower chance of getting cancer (say 25%). Consider the sentence "I will get lung cancer." And consider the state of affairs "I am a smoker". As it turns out the latter causes both the obtaining and non-obtaining of the former. If you live a long cancer-free life, you have the smoking to thank (for it suppressed the disorder), but if you end up with cancer, then your smoking will figure centrally in that explanation (the disorder caused the smoking in a round-about way, but the cancer is almost certainly due to smoking and not the disorder if you smoked regularly).[35] The important upshot here is that, for Zardini, failures of contraction are explained by a kind of instability that sentences have and this instability can naturally be understood as a sort of metaphysical correlate of ambiguity or underdetermination. The main point I want to make is that Zardini denies contraction by denying **P1:** NO EQUIVOCATION.[36]

Zardini is not alone in appealing to phenomena of this sort. Caret and Weber (2015) explicitly invoke ambiguity in explaining failures of contraction:

> "It is thought that sentences express propositions, and one either rejects, or does not reject, a proposition; on previous conceptions, one sentence can have only one proposition as its content. [...] The VCurry sentence C does indeed appear to express two different propositions [...] On the one hand, C expresses that C is true; and on the other, it expresses that C is absurd. Prima facie, these are two different propositions, and at this stage of our understanding about such pathologies, it seems prudent to take the content of C at fact[sic] value. When the boy cries wolf a second time, his utterance is the same token, but a different type. Upon repetition, it comes to mean something else.
> The salient point is not merely that the sentence C expresses two things. Lots of sentences do that. But in most cases we expect that there is a singular proposition expressed by a

[35] To borrow a famous example, if we think it is simply metaphysically indeterminate whether "There will be a sea battle tomorrow", then we need to only find the existence of a state of affair that causes *both* the obtaining and the non-obtaining of this state of affairs. If there's some state of affairs floating around right now that would figure centrally in future histories of why there was (or was not) a sea battle. If there's one that figures in both, then we have such a state of affairs.

[36] I do not discuss in detail a recent and compelling account by Rosenblatt (2019), according to which violations of contraction may arise where a sentence is equivalent*—a technical notion of Rosenblatt's meaning: intersubstitutible preserving validity—to its own negation. This notion is given flesh by invoking the bilateralism of Restall (2005); Ripley (2017). Rosenblatt thinks that such sentences are extremely rare (i.e. only occur for paradoxical sentences). I think that Rosenblatt has in mind a similar intuition to many of the other accounts I discuss: one sentence expresses two distinct propositions: a sentence and its own negation; however, it is unclear to me whether he would explicitly reject NO EQUIVOCATION. First, his own wording is that a sentence is *equivalent* to its own negation. And equivalence of two sentences is unproblematic. If it turns out that Rosenblatt is safe from the sort of objection that I advance, then I welcome him onboard as an ally: he too sees that content may be non-contractive owning to its *multi-facetedness*. In this case, my critique of him concerns the explanation of why content may fail to contract. See below for further detail.

> given sentence—the conjunction of its contents—that includes every proposition it expresses as a part. Matters are not so simple with the VCurry." (Caret and Weber, 2015, pp.66–7)

A sentence is non-contractive because it expresses *two propositions*. This is, of course, an explicit denial of **P1:** NO EQUIVOCATION. It is not news that logics without contraction can be used for equivocation (see e.g. Lewis (1982)). They have long been used for making sense of vague phenomena in the style of Łukasiewicz many-valued logic (see also (Slaney, 2010)).

By now I hope that the evidence is clear and overwhelming that most take for granted the connection between **P1:** NO EQUIVOCATION and **C:** CONTRACTION. Those who think giving up contraction is non-sensical, can't understand how one can give an account of reasoning with sentences that have definite senses absent contraction; those who give up contraction try to argue for the plausibility of a sentence being underdetermined or ambiguous in some way or other. I'd like to briefly consider one other class of accounts which see violations as a more general phenomena. The primary reason for thinking this has to do with thinking of consequence in terms of a kind of "resource management".

Girard (1995) distinguishes two senses of "and"—or of the comma—one of which allows for contraction and one of which does not. For example, suppose that you have \$1 and that a pack of cigarettes costs \$1. Then if you have \$1 it seems to follow that "You may buy a pack of Camels" and "You may buy a pack of Marlboros". But it does not follow that "You may buy a pack of Camels and a pack of Marlboros". There thus seems to be a distinction between two senses of "and", one of which entails that each may be constructed, and one which entails that both may be constructed. A similarly ambiguity appears on the left as well, for there are two senses of "and" which allow us to make sense of how we could derive *both* of the above:

"I have \$1" and "I have \$1" $\vdash$ "I may buy a pack of Camels" and "I may buy a pack of Marlboros",

but another sense of "and" where the two above contract:

"I have \$1" and "I have \$1" $\nvdash$ "I may buy a pack of Camels" and "I may buy a pack of Marlboros", It is this *ambiguity* between the two senses of "and" or between the use of the commas on the left- and right-hand sides which explains why the possibility of violations of contraction have been overlooked. If we think of sentences as resources with which to reason (or with which to build proofs) rather than as things which have truth-values (such that two occurrences have the same content as one occurrence), then we open up the possibility of violations of contraction.

The reason that this involves a denial of **P1:** No Equivocation involves how the "contractive" conjunction behaves. Consider a classical proof of: $A \wedge (A \to B) \vdash B$. There is a sense of $\wedge$ under which this works (multiplicative), but there's another sense under which it doesn't. Namely a sense under which $A \wedge (A \to B)$ is indeterminate between A and $A \to B$. It can express A and it can express $A \to B$, but can't really express both. This involves a kind of ambiguity. This is the same sort of ambiguity we saw concerning "I may buy a pack of Camels" and "I may buy a pack of Marlboros" above. Their conjunction expresses each of them, but not both at the same time. This sense of $\wedge$ (the additive one) allows the derivation with two copies of the sentence: $A \wedge (A \to B), A \wedge (A \to B) \vdash B$, but not with a single copy. Thus, two copies of a sentence gets you something one copy does not precisely because the sentence is indeterminate between two different contents. This is precisely a denial of **P1:** No Equivocation.[37]

3.2.1.3 What is and is not multi-facetedness

So far I've argued that considerations for and against contraction all seem to share an assumption. They all assume that if in a proper accounting of reasoning the content of sentences are all properly disambiguated, then contraction must hold. The antecedent of this I've called **P1:** No Equivocation. But the link between No Equivocation and Contraction is only secured by **P2:** Sameness is Redundant. While this is treated as "obvious", no argument is given for it, and I intend to deny it. Before I fill in my positive

[37] This insight has been used e.g. by Mares and Paoli (2014) to argue for a non-contractive solution to paradox. Traditionally logics have run together internal consequence and external consequence. So just as conjunction is ambiguous (as above) so too is the conditional. This gives us an understanding of how two copies of Curry yield any falsehood, but one copy doesn't.

This has also been used in non-logical contexts. For example, Barker (2010) has argued that we should understand many deontic modals as involving resource management of e.g. actions.

account I want to try to better characterize what I'll call the multi-facetedness of content. While I think that vague, ambiguous, or otherwise underdetermined sentences *are* multi-faceted, that they are such is not strictly speaking a feature of their being multi-faceted—i.e. this goes beyond multi-facetedness. The main distinction I want to focus on here is that the *facets* in virtue of which a sentence is *multi*-faceted aren't complete senses or propositions. What is problematic about vague, ambiguous, or underdetermined sentences is that a single sentence may express multiple, distinct propositions—each of which is capable of standing on its own.

Let's start with what multi-facetedness is not.[38] It may be that a single sentence expresses distinct contents. This may happen in perfectly unharmful ways. For example post-Kaplan (1979), the following is taken to express distinct contents:

(Peter) Peter put it in the drawer.

such that each assertion of this sentence may mean something different (the sentence considered abstractly, therefore, underdetermined with respect to them), but each of those assertions share a single *character* (hence their homophony). It seems plausible that the character of Peter implies e.g. "It is in the drawer", but strictly speaking this implication only holds when the distinct occurrences of "it" are lined up in the appropriate way and only because of the relation between those contents.

We say that the sentence "Peter put it in the drawer" is therefore, context sensitive. It has a single "meaning" (namely: its character), but it is capable of expressing *distinct contents*. Thus, two occurences of *the same sentence* may express distinct contents. This is not what I mean by multi-faceted.

Further, a sentence such as:

(Paul) Paul put it in the basket because it was wet.

Is under-determined along a further dimension. Different utterances of this noise vary not only with respect to content, but with respect to character as well. In particular, it makes

[38]I explain that multi-facetedness is not context-sensitivty, ambiguity or polysemy, vagueness, or any other sort of semantic *under*determination. I don't mean to rule out an explanation of these phenomena *in terms of* multi-facetedness (I think such an analysis could be promising). Rather, I mean to explain that a sentence may be multi-faceted without being any of those things, and that the possibility of multi-facetedness does not require allowing any form of semantic underdetermination.

a difference whether we require that the two instances of *it* co-refer or not. These are two different grammatical structures that the sentence can have:[39]

1. Paul put it_1 in the basket because it_2 was wet.
2. Paul put it_1 in the basket because it_1 was wet.

For example, perhaps Paul put a wet-shirt in the basket because the shirt was wet. Or perhaps he put a towel in the basket because the basket was wet. The important thing to note here is that even once we make the content/character distinction we still have two ways of understanding this sentence. Particular assertions won't have a single content associated with them.[40] Such a sentence not only expresses distinct contents, but distinct characters as well. We can therefore understand it not only context-sensitive, but ambiguous or under-determined.

The above won't do for the sort of account I intend to give. Nevertheless, I think it helps get the idea of a *facet* onto the table. A facet is an aspect of a sentence's meaning. Something that the sentence contains, but which is not, strictly speaking, an autonomous proposition. If a single sentence could express two distinct contents, we would call such a sentence ambiguous. If a single sentence contains multiple facets, then all we mean to observe is that it is involved in partially distinct patterns of use. Such a phenomenon could be potentially widespread, but I limit myself here to carefully constructed cases.[41]

Let me give an example to illustrate how it works. I stick to examples of this shape, involving something like analyticity, because they seem to be the easiest such cases. Analyticity

[39] It has been pointed out to me that there's actually a third variant where the second "it" is non-referring as in "it was wet" (e.g. because it was raining).

[40] Similar problems can be constructed using quantifiers. For example.

(Mary) Mary went to the library to look for a book.

Has two readings depending upon how we bind the quantifier "a" before book. That is, either:

1. There is a book such that Mary went to the library to look for *it*.
2. There is a relation between people and books—the relation of looking for—and the extension defined by Mary taking the first argument place is non-empty.

Now, because of these sorts of ambiguities, we can imagine situations where two instances of each of these appear and could take divergent syntactic readings. If both occur in a single context it could happen that two occurrences imply something that one occurrence does not, hence a violation of contraction. But if all violations of contraction are like this, then we must deny No Equivocation.

[41] Suppose that there are contexts in which two sentences are intersubstitutible (though this intersubstitutibility need not hold globally). We can call this usage of these sentences a *facet*.

involves a kind of containment that involves the same sort of intuitions as multi-facetedness. I do not think (nor do I want to claim) that all examples involve analyticity, however. The general shape of the example is as follows. First, we want one sentence to be a facet of another sentence (suppose that A is a facet of B). This means that the following holds:

$$\frac{A, \Gamma \vdash \Delta}{B, \Gamma \vdash \Delta}$$

Next, we have to exhibit the following pattern of implications:

$$B, B, \Gamma \vdash \Delta$$
$$B, \Gamma \nvdash \Delta.$$

That is the top sequent of contraction holds and the bottom doesn't, so the rule of contraction is violated. To do this, we give an intuitive example where:

$$B, \Gamma \nvdash \Delta,$$
$$B, A, \Gamma \vdash \Delta.$$

And in light of the second one, we have (via the facet relation) that $B, B, \Gamma \vdash \Delta$. It is important to keep in mind that the facet relation (in light of the former) requires $A, \Gamma \nvdash \Delta$ (since we can read sequent rules contra-positively, i.e. if the bottom fails to obtain, so does the top). I mention this to explain all of the relevant constraints in play.

Here's the example. It is analytically contained with the concept of lying that one intends to deceive. That is, suppose:

L = "Louise told Ron a lie"

D = "Louise intended to deceive Ron"

Then, to say that D is a facet of L is to say the following holds:

$$\frac{D, \Gamma \vdash \Delta}{L, \Gamma \vdash \Delta}$$

Next, let W = "Louise wronged Ron" and let us suppose that $L \vdash W$, i.e. telling a lie is to harm someone. Sometimes—i.e. in light of further considerations—lying is not only permissible but does not wrong the recipient. For example, if B = "If Louise told the truth, it would hurt Ron considerably", then:

$$L, B \nvdash W.$$

In light of the fact that telling the truth would cause Ron harm, it does not follow that Louise wrongs Ron by lying. Perhaps lying is even the thing to do. But of course, if the lie is primarily done out of an intent to deceive, then Ron will be wronged. If it is primarily done out of a motive to prevent harm, then it could be permissible (even though the intention to deceive must, in some sense, be active. What I mean is that it seems clear:

$$L, B, D \vdash W.$$

But then (since D is a facet of L), that $L, L, B \vdash W$, and so together we have:

$$\begin{aligned} L, L, B &\vdash W \\ L, B &\nvdash W \end{aligned}$$

To help see this, it may help to arrange the information into a particular presentation. I don't want to claim that $\vdash$ has anything to do with such arrangements: it is meant to represent a relation of reasoning; but such arrangements can make such reasoning perspicuous:[42]

> Louise lied to Ron. But had she told the truth she would have caused considerable harm.
> True, but Louise's intention was to deceive Ron, and so she wronged him.

[42] An alternative would be to say that although the implication with two occurrences of the same sentence is intuitively "strange", several commitments we have made suggest we ought to accept it in our formal representation of the consequence relation. Dicher (2020) for example argues that while multiple conclusions are quite strange in ordinary practices of reasoning, many of our commitments concerning reasoning give us reason to admit them into our formal representation of that reasoning. [...] suggested this alternative to me.

While I've added some flourishes to the information, I hope that the reader can recognize them as pragmatic fluff: they help the argument digester understand how the argument works, but are strictly speaking not part of the semantic contents of the various sentences.

Second: it might be that one has particular quibbles with this example (or any other example). I only want to claim it's plausible that multi-facetedness is a feature of content in a wide variety of cases. But to add some plausibility in this case, it's important I think to understand that L doesn't express D, but that D is a facet of a particular aspect of L's usage (namely as a premise in the argument). But there may be premises which justify us in concluding that D that do not justify us in concluding L:

$$\Gamma \vdash \Delta, D \quad \text{but} \quad \Gamma \nvdash \Delta, L.$$

For example if it is clear that Louise intended to deceive Ron, but perhaps what she said happens to be true (and so is not a lie).

Addressing the specifics: while $L, B \nvdash W$ suggests that telling the truth would not wrong Ron, we needn't understand things in this way. It may be that the thing to do for Louise is to not say anything. If telling the truth would cause Ron harm, but she has a strong motive to deceive him (perhaps lying carries some benefit to her), then staying silent may be the thing to do.

Summing up: I've claimed that logics concerned with contraction think that its denial is incompatible with **P1:** NO EQUIVOCATION. I think that this is a mistake and that a suppressed premise **P2:** SAMENESS IS REDUNDANT is needed for the argument. If we could deny **P2** instead, then we would open room for a non-contractive logic that doesn't rely on e.g. equivocation, ambiguity, etc. to make this denial palatable. I claim that the feature of content which allows for this is its *multi-facetedness*. I hope to have made clear so far how it is at least plausible that one and the same content may be involved in a good implication where it has two occurrences, but not where it has one occurrence. Of course, I've claimed that a sentence like L above is *not* ambiguous. But on a standard truth-conditional account, its truth conditions would be intimately tied to the truth conditions of D. But recall that this is precisely what I aim to deny. I've started with the notion of good implication as basic and used that to furnish content (rather than building a notion of implication out of content). In the next section I'll use this account to make the account of multi-facetedness more precise.

3.2.2 Multi-Facetedness and Inferential Productivity

With my account of meaning in terms of contribution to good implication in place, we can see what it means for contraction to fail, and what it means for a sentence to be multi-faceted. The idea that a sentence is multi-faceted is captured by the idea that a sentence may be involved in distinct sets of implications. For example, it may be that A (as a premise) makes contribution Z to good implication as a premise, but that $X, Y \subsetneq Z$ are each themselves proper contributions—where "proper" here means they meet Constraint Two: they fix an equivalence class (see p. 15 above). That is: A has among its contribution so-to-speak sub-contributions. In a very strict sense: Z contains more information than all of the sub-regions contain. This is because—among other things—how the various facets of A compose is not given by how they behave independently of one another. Though that information is contained in Z. To make the point as clear as possible, we may have:

$$
\begin{array}{ccc}
A, \Gamma_1 \vdash \Theta_1 & B, \Delta_1 \vdash \Lambda_1 & C, \Pi_1 \vdash \Xi_1 \\
\vdots & \vdots & \vdots \\
A, \Gamma_n \vdash \Theta_n & B, \Delta_m \vdash \Lambda_m & C, \Pi_l \vdash \Xi_l.
\end{array}
$$

with

$$
\begin{aligned}
X &= \{\langle \Delta_i, \Lambda_i \rangle | 1 \leq i \leq m\}. \\
Y &= \{\langle \Pi_i, \Xi_i \rangle | 1 \leq i \leq l\}. \\
Z &= \{\langle \Gamma_i, \Theta_i \rangle | 1 \leq i \leq n\} \\
&\supsetneq X \cup Y.
\end{aligned}
$$

Whether or not X and Y are proper inferential roles will depend upon what the entire consequence relation looks like,[43] but assuming they are proper inferential roles, we call them facets. Importantly facets are not sentences that another sentence (in this case A) expresses. While A may contain these facets as a premise, its relation to e.g. B and C might not be so straightforward (depending upon how A, B, and C behave as conclusions in this case).

Let's explore a concrete example, drawing on some of the material from earlier.[44]

Consider the following sentences:

M**:** When Jones shot Jack, it was murder.

U**:** Jones had no justification for shooting Jack.

K**:** Jones killed Jack by shooting him.

Each of these sentences are unambiguous and distinct (or at least we should be able to read them as such). Further, most would understand U and K to be analytically contained in the first M, since "Murder is unjustified killing" is often accepted as analytic. Now, at least as premises, anything that follows from U or follows from K should also follow from M:

$$\frac{U, \Gamma \vdash \Delta}{M, \Gamma \vdash \Delta} \qquad\qquad \frac{K, \Gamma \vdash \Delta}{M, \Gamma \vdash \Delta}$$

Clearly from everything I've said thus far, U and K are *facets* of M. I think we can imagine cases where $M, M, \Pi, \Gamma \vdash \Delta$ but $M, \Pi, \Gamma \not\vdash \Delta$. For example suppose:

D**:** Jones will receive a prison sentence.

V**:** Jack was acting erratically and threatening violence.

[43]The easiest way to secure that they are, is to postulate premissory (or conclussory) roles that behave exactly as they do in exactly those regions. That is: postulate a B and C which are involved in exactly those implications. Then X and Y will be proper inferential roles. Of course if such X and Y exist, then there would be a "disambiguation" of their facets, namely B and C, so such postulation, while sufficient, ought not (and is not) necessary to get the desired result. Even so, such disambiguation is at the level of *role* not at the level of sentence. We shouldn't think of A as expressing B or C.

[44]I've set up enough to motivate understanding the lying case from earlier in these terms. I invoke a similar example here. There's a reason for this. The law is full of definitions (so analytic relations between contents), but the law is also what H.L.A. Hart (1948) has called "open textured," by which he means that legal concepts often go beyond the definitions given of them.

It seems plausible that if Jones killed Jack, he'll receive a prison sentence, certainly it's true if he murdered him (let Γ and Δ be the relevant contextual information need to fill in the details, i.e. $K, \Gamma \vdash \Delta, D$ and thus $M, \Gamma \vdash \Delta, D$).

But if we add V, then this will not follow, i.e. $K, V, \Gamma \not\vdash \Delta, D$.[45] But if we add U—i.e. the explicit premise that the killing was unjustified—then we get the implication back: $K, V, U, \Gamma \vdash \Delta, D$. Because of the facet-relation we thus have:

$$M, M, V, \Gamma \vdash \Delta, D$$

As I mentioned earlier, to hear this example "intuitively" might require adding some pragmatic information or arranging the premises into a narrative:

> Jack murdered Jones. No, Jones was acting erratically and threatening violence. The facts don't lie: Jones murdered him, and will, therefore, receive a prison sentence.

My claim (again) is that the content of this narrative is what the implication is based upon. But that the implication is good depends solely on the content (not on the narrative structure).

Now, in these cases, in addition to multi-facetedness a second feature was needed; namely that the facets be *inferentially productive*. For it might be the case that a sentence contains multiple facets which fail to produce a violation of contraction. For example, perhaps "Jim is a human" contains analytically "Jim is an animal" as a facet, but this need not entail any violation of contraction. What is needed in addition is that the facet be inferentially productive with respect to the sentence. For example, in the previous example, U and K turned out to be inferentially productive with respect to one another. What this means is that there are good implications that they together contribute towards that neither one singly contributes toward. Importantly: U is inferentially productive with respect to M (U together with M contribute to good implications that M does not contribute to).

The idea of inferential productivity of course is quite intelligible apart from contraction violations. For example, it is generally accepted as a schema that A and $A \to B$ are

[45] It might be harder to see $M, V, \Gamma \not\vdash \Delta, D$. Technically our premise set contains contradictory information: he murdered, but the victim was acting erratically—so he wasn't justified! We can alter the set-up in minor ways if needed. Perhaps we strengthen V—he wasn't justified!, or we observe that even if it was still a killing, Jack's behavior is likely to lead to more lenient sentencing (Jones had partial but incomplete justification).

inferentially productive with respect to one another (neither implies B, but together they contribute to the goodness of an implication where B figures as a conclusion, namely: $A, A \to B \vdash B$).

Let me briefly summarize the philosophical account so far introduced. I am advocating understanding meaning in terms of contribution to good implication. On such an account the meaning of a sentence does not vary in the various contexts in which it appears, and crucially does not vary even where two occurrences render an implication good, but one would not. To bring this point out, I explained that violations of contraction occur where a sentence has a facet that is inferentially productive. Crucially facets are not sentences, and there need be no sentence which (nor even particular premissory/conclussory contribution of a sentence) we can identify a facet as. There is no equivocation involved in identifying facets of sentences. Rather, we should understand facets as patterns of use that a sentence may be involved in. It may happen that a facet is inferentially productive with respect to the very contribution of which it is a facet. In such a case, we have a violation of contraction.

As it turns out, the notions I've introduced so far can be made precise in the formal semantics introduced in the previous chapter. I shall now introduce the formal details.

3.2.3 A Precise Account of Non-Contractive Content

According to my account violations of contraction happen when two features of inference come together. First a sentence must be multi-faceted. This means that there is some pattern of its use in implication, which is not equivalent to it. The second is that this *facet* of a sentence be *inferentially productive*. We call an inferential role productive (typically *with respect to some other role*) if its addition makes a strictly larger contribution to good implication. Both of this notions are made precise here. In the section that follows I illustrate how they manifest in some familiar examples.

Definition 3.2.1 (Facet)**.** Let $\langle X, Y \rangle$ be a proper inferential role. Call $\langle U, V \rangle$ a **facet** of $\langle X, Y \rangle$ iff $\langle U, V \rangle$ is a proper inferential role and $U \subseteq X$ or $V \subseteq Y$. Often, we will simply refer to U or V as facets of $\langle X, Y \rangle$ (as a premise/conclusion) where lack of ambiguity allows.

Intuitively, a facet of an inferential role is a particular way of thinking about that role, a characterization of it. Perhaps the easiest way to understand this is in terms of a *containment* relation between roles (of the sort I just described). For example, suppose that $[\![B]\!]_P \subseteq [\![A]\!]_P$. Then $[\![B]\!]_P$ characterizes a facet of A. This means that whenever B occurs as a premise, we could replace it with A:

$$\frac{B, \Gamma \vdash \Delta}{A, \Gamma \vdash \Delta}$$

NB: it is important to keep in mind that B is not a facet of A here. It may be that $[\![B]\!]_C \not\subseteq [\![A]\!]_C$, and in fact that the conclussory behavior of these diverge. Further, A may have a facet which is not characterized by any sentence. Both of these are important to note because my thesis is that facets are not expressed contents. That a sentence is multi-faceted is not to be understood in terms of a sentence being ambiguous or anything of that sort.

Fact 3.2.2. Every non-empty, proper inferential role has at least two facets (namely itself and the empty facet).[46]

Definition 3.2.3 (Inferential Productivity)**.** Let X and Y be closed. We say that Y is productive with respect to X iff:

$$X \not\supseteq (X^{\curlyvee} \sqcup Y^{\curlyvee})^{\curlyvee}.$$

The paradigmatic case is where $X \subsetneq (X^{\curlyvee} \sqcup Y^{\curlyvee})^{\curlyvee}$. Strictly speaking, the definition covers the case where there occurs something in the latter which does not occur in the former. The less strict definition is intended to account for a wide variety of cases that may arise in a substructural setting. If we restrict our attention to monotonic $\vdash$, then we get the definition in terms of '$\subsetneq$'.[47]

Inferential productivity is, to be precise, something that holds between *contributions* and not between sentences. However, I sometimes talk as if it holds between sentences. This is

[46] If $\langle X, Y\rangle$ is a PIR with both X and Y non-empty, then $\langle X, Y\rangle$ has four facets.

[47] I don't prove this, but here's a quick sketch of why this is true. Suppose $X = \langle\{A\}, \emptyset\rangle^{\curlyvee}$ and $Y = \langle\{B\}, \emptyset\rangle^{\curlyvee}$.The following is an admissible rule given monotonicity:

$$\frac{A, \Gamma \vdash \Delta}{A, B, \Gamma \vdash \Delta}$$

So whenever $\langle \Gamma, \Delta\rangle \in X$, it must be in $(X^{\curlyvee} \sqcup Y^{\curlyvee})^{\curlyvee}$. But the converse doesn't hold since it is possible that $A, B, \Gamma \vdash \Delta$ but $A, \Gamma \not\vdash \Delta$.

particularly unproblematic if, for example, both the premissory and conclussory roles of B are productive with respect to the premissory and conclussory roles, respectively, of A.

There are some very general things we can say about inferential productivity. Suppose that our consequence relation satisfies the schema $\Gamma, p \vdash p, \Delta$ (thus $\vdash$ is supra-classical). Then we know some basic facts about inferential productivity.

Proposition 3.2.4. Suppose that $\vdash$ satisfies the schema: $\Gamma, p \vdash p, \Delta$. Then, supposing that p is neither tautologous nor contradictory, a sentence's premissory and conclussory roles are always productive with respect to one another. That is:

$$[\![p]\!]_P \not\supseteq ([\![p]\!]_P^{\curlyvee} \sqcup [\![p]\!]_C^{\curlyvee})^{\curlyvee}.$$

The same holds if we exchange all the $[\![\cdot]\!]_P$ and $[\![\cdot]\!]_C$.

Proof. Consider that $\langle\{p\},\{p\}\rangle^{\curlyvee} = \mathbf{P}$. Reminder that $\mathbf{P}$ is the set of all points in inferential space. So the claim that $\langle\{p\},\{p\}\rangle^{\curlyvee} = \mathbf{P}$ is equivalent to the claim that for arbitrary $\langle\Lambda, \Delta\rangle \in \mathbf{P}$ we have: $\Lambda, p \vdash p, \Delta$. ■

We have to rule out tautologies and contractions because, for example *nothing* is productive for $[\![A \vee \neg A]\!]_C$ (since its contribution is already $\mathbf{P}$).

If $\vdash$ satisfies this schema[48] there are a number of instances of inferential productivity downstream from this. I emphasize they are downstream because the inferential productivity flows from the schema $\Gamma, p \vdash p, \Delta$ and e.g. the behavior of the negation and the conditional are parasitic on this. There is nothing inherent to them that causes inferential productivity.

Fact 3.2.5. The negation of a sentence is always productive with respect to that sentence. In fact it is maximally productive. This fact amounts to the observation that:

$$\langle\{A, \neg A\}, \emptyset\rangle^{\curlyvee} = \mathbf{P}.$$

Fact 3.2.6. For any sentence A, $[\![A \to \phi]\!]_P$ is inferentially productive with respect to $[\![A]\!]_P$. This is because of the valid implication schema of modus ponens.

Theorem 3.2.7. *Say that a sentence A is involved in a failure of contraction if there exists Γ and Δ such that one of the following does* not *hold:*

[48] By this schema I mean CO. Recall that this the schema: $\forall\langle\Lambda, \Delta\rangle \in \mathbf{P} \forall A \in \mathcal{L}(\Lambda, A \vdash A, \Delta))$

$$\frac{A, A, \Gamma \vdash \Delta}{A, \Gamma \vdash \Delta} \qquad\qquad \frac{\Gamma \vdash \Delta, A, A}{\Gamma \vdash \Delta, A}$$

We have that A is involved in a failure of contraction iff A has an inferentially productive facet.

Proof. ($\Rightarrow$) Suppose that A is involved in a failure of contraction. Without loss of generality, let it be $A, A, \Gamma \vdash \Delta$ but $A, \Gamma \nvdash \Delta$. Clearly $[\![A]\!]_P \subseteq [\![A]\!]_P$ and so we have a facet of A. And clearly this facet is productive since $\langle \Gamma, \Delta \rangle \in ([\![A]\!]_P^{\curlyvee} \sqcup [\![A]\!]_P^{\curlyvee})^{\curlyvee}$ but $\langle \Gamma, \Delta \rangle \notin [\![A]\!]_P$.[49]

($\Leftarrow$) Suppose that A has a facet which is inferentially productive. Without loss of generality suppose that $X^{\curlyvee\curlyvee} \subseteq [\![A]\!]_P$, and suppose X is inferentially productive. Then let $\langle \Gamma, \Delta \rangle \in ([\![A]\!]_P^{\curlyvee} \sqcup [\![X]\!]^{\curlyvee})^{\curlyvee} \setminus [\![A]\!]_P$.Then for arbitrary $\langle \Pi, \Lambda \rangle \in X^{\curlyvee\curlyvee}$ we have:

$$A, \Pi, \Gamma \vdash \Delta, \Lambda,$$

but $A, \Gamma \nvdash \Delta$. By supposition (since $X^{\curlyvee\curlyvee}$ is a facet of A) we have:

$$A, A, \Gamma \vdash \Delta,$$

and hence a violation of contraction. ■

[49]This might seem like my definitions aren't doing much philosophical work. If this is what a violation of contraction is, then haven't I simply said that a sentence is non-contractive, provided it is non-contractive? I don't think this objection is quite right. I am being maximally tolerant concerning $\vdash$ here. This means in the extreme case that A's premissory role might be productive with respect to itself. For example, consider a consequence relation consisting of exactly one element: $\{\langle \{A, A\} \cup \Gamma, \Delta \rangle\} = \mathbb{I}$.

The proof of ($\Leftarrow$) in any case shows why this isn't a cheap trick. That is, that more interesting cases of multi-facetedness are doing real work.

3.2.3.1 Paradoxical Sentences

I have argued that multi-facetedness is a unifying explanation for failures of contraction and vindicated that thought in a formal framework. But where does that leave the paradoxes? The original motivation for giving up contraction was to find a solution to Curry; later this was used to find a solution to the liar as well. Can I say anything interesting about those in particular? It's actually easy to see them as having a common form given the machinery here. First let's consider the liar:

$$\lambda =_{df.} \neg Tr(\lambda).$$

We may reason:

$$\begin{aligned}
\langle [\![\lambda]\!]_P, [\![\lambda]\!]_C \rangle &= [\![\lambda]\!] \\
&= [\![\neg Tr(\lambda)]\!] \\
&= \langle [\![\neg Tr(\lambda)]\!]_P, [\![\neg Tr(\lambda)]\!]_C \rangle \\
&= \langle [\![Tr(\lambda)]\!]_C, [\![Tr(\lambda)]\!]_P \rangle \\
&= \langle [\![\lambda]\!]_C, [\![\lambda]\!]_P \rangle.
\end{aligned}$$

The result of this is that

$$\langle \{\lambda\}, \emptyset \rangle^{\curlyvee} = \langle \emptyset, \{\lambda\} \rangle^{\curlyvee}.$$

That is: any context $\langle \Gamma, \Delta \rangle$ is such that λ's adjunction qua premise makes a good implication iff its adjunction qua conclusion does. So we notice something immediate about the liar. It contains a facet, namely its own conclussory role, and that facet is productive (assuming we are working in a classical setting, a sentence's premissory and conclussory roles are always inferentially productive with respect to one another).

Proposition 3.2.8. The liar has an inferentially productive facet.

I'd like to highlight a few very important features of how the liar has been handled here. First, I want to emphasize the claim:

The liar is non-contractive because it has a facet with which is inferentially productive.

I want to emphasize this in order to highlight why the liar is paradoxical on my telling. We produced an interesting facet of the liar from its definition alone (together with the definition of '$\neg$'). I think this has consequences for how we should understand the order of explanation. The reason that the liar doesn't contract isn't because it contains a facet that is its negation. Rather, the behavior of the negation is what explains the fact that it contains that facet. It is inferentially productive because a sentence's conclussory role is always productive with respect to its premissory role (so long as the schema $\Gamma, p \vdash p, \Delta$ holds). Just to be perfectly clear on what I mean with the previous sentence: negation explains why the liar has a particular facet (namely it explains why $[\![\lambda]\!]_P = [\![\lambda]\!]_C$), but this doesn't on its own explain why the liar is paradoxical or why contraction is violated. For that we also need to observe the inferential productivity, which is a fact that we don't need to invoke negation to explain.

This can be brought more clearly into focus by considering how Curry works.[50] Suppose we have a sentence:

$$\kappa_q = Tr(\kappa_q) \to q.$$

[50] I think it makes a substantive technical as well as philosophical difference whether one speaks of Curry as a single sentence ($\kappa = Tr(\kappa) \to \perp$) or as a schema (i.e. $\kappa_q = Tr(\kappa) \to q$). This sentence is only problematic if q is a contradiction or falsehood. So it is really the schema that is problematic. The '$\perp$' presentation is in a way equivalent to the liar given certain presuppositions—this is unsurprising since negation can often be defined as $\neg p =_{df.} p \to \perp$. If we understand $\perp$ as:

$$\frac{}{\perp, \Gamma \vdash \Delta} \qquad\qquad \frac{\Gamma \vdash \Delta}{\Gamma \vdash \Delta, \perp}$$

i.e. such that $[\![\perp]\!] = \langle \mathbf{P}, \mathbb{I} \rangle$, then it's easy to show that κ behaves like λ, i.e. $[\![\kappa]\!]_P =_{df.} [\![\kappa]\!]_C$:

$$\begin{aligned}
\langle [\![\kappa]\!]_P, [\![\kappa]\!]_C \rangle &= [\![\kappa]\!] \\
&= [\![Tr(\kappa) \to \perp]\!] \\
&= \langle [\![Tr(\kappa) \to \perp]\!]_P, [\![Tr(\kappa) \to \perp]\!]_C \rangle \\
&= \langle [\![Tr(\kappa)]\!]_C \cap [\![\perp]\!]_P, (([\![Tr(\kappa)]\!]_P)^{\curlyvee} \sqcup ([\![\perp]\!]_C)^{\curlyvee})^{\curlyvee} \rangle \\
&= \langle [\![\kappa]\!]_C \cap \mathbf{P}, ([\![\kappa]\!]_P^{\curlyvee} \sqcup \mathbb{I}^{\curlyvee})^{\curlyvee} \rangle \\
&= \langle [\![\kappa]\!]_C, ([\![\kappa]\!]_P^{\curlyvee} \sqcup \langle \emptyset, \emptyset \rangle)^{\curlyvee} \rangle \\
&= \langle [\![\kappa]\!]_C, [\![\kappa]\!]_P^{\curlyvee\curlyvee} \rangle \\
&= \langle [\![\kappa]\!]_C, [\![\kappa]\!]_P \rangle
\end{aligned}$$

I think we need to consider the schema Curry separately, however, because explanations of $\perp$-Curry and q-Curry often come apart.

It is easy to see that $[\![\kappa_q]\!]_P \subseteq [\![\kappa_q]\!]_C$:

$$\begin{aligned}
[\![\kappa_q]\!]_P &= [\![Tr(\kappa_q) \to q]\!]_P \\
&= [\![Tr(\kappa_q)]\!]_C \cap [\![q]\!]_P \\
&= [\![\kappa_q]\!]_C \cap [\![q]\!]_P \\
&\subseteq [\![\kappa_q]\!]_C.
\end{aligned}$$

Actually, if we assume that q is neither a tautology nor a contradiction (for simplicity, assume it is just an atomic sentence), then we have that $[\![\kappa_q]\!]_P \subsetneq [\![\kappa_q]\!]_C$. In short, $\langle \{q\}, \emptyset \rangle$ is in the latter, but not the former (we can show $q \vdash \kappa_q$ as an instance of $\Gamma, q \vdash q, \Delta$, but this does not hold for $\kappa_q, q \vdash$, which requires $q, q \vdash$, which we do not have).

Nevertheless, premissory and conclussory roles are still productive with respect to one another (if the schema $\Gamma, p \vdash p, \Delta$ holds—in fact, for the case of κ_q, the schema $\Gamma, p, p \to r \vdash r, \Delta$ sheds more light).[51]

Proposition 3.2.9. κ_q contains a facet with which it is inferentially productive.

Proof. The construction of the facet is given above. In particular $[\![\kappa_q]\!]_P$ is a facet of $[\![\kappa_q]\!]_C$. It is straightforward to show that: $\vdash \kappa_q, \kappa_q$ (since $\kappa_q \vdash \kappa_q$) but $\nvdash \kappa_q$. ■

3.2.3.2 Relation to Other Substructural Solutions to Paradox

I'd like to close my discussion of paradox by exhibiting the relationship my own account has to other substructural solutions to paradox. My hope isn't necessarily to demonstrate that my account is right and theirs wrong, but that my account is at least more explanatory. I showed that sentences fail to contract if they contain a facet which is inferentially productive, and I showed that paradoxes have precisely these two features. What is distinctive about

[51] The violation of contraction related to q-Curry is contraction of conclusions. I set it up this way because these facts concerning q-Curry's multi-facetedness follow straightaway from its definition and don't require *any* further assumptions. If we assume a classical set-up (i.e. assume the schema $\Gamma, p \vdash p, \Delta$), then $\langle \{\kappa_q\}, \{q\} \rangle^{\curlyvee}$ will be a facet of $[\![\kappa_q]\!]_P$, which allows us to show how this manifests in κ_q's behavior as a premise. But the existence of this facet relies on the classical shape of the consequence relation. Just to underline this claim explicitly: the multi-facetedness (related to its premissory role) follows immediately from mere definitions and requires *no* assumptions about the shape of the consequence relation (not even the assumption of CO).

paradoxes is typically the facets in question. Often it is the conclussory role of the paradoxical sentence that is a facet of its own premissory role (or vice-versa). The resources of self-reference combined with e.g. a naive theory of truth seem needed to generate this. That the conclussory role of a sentence is inferentially productive of its premissory role is not strange behavior, however. For, we typically think that all non-tautologous/contradictory sentences behave this way. To be clear, all sentences seem to involve the following:

$$[\![p]\!]_P \subsetneq \langle\{p\},\{p\}\rangle^{\curlyvee} = \mathbf{P}.$$

This says that $\langle\{p\},\{p\}\rangle$ makes a *maximal* contributon to good implication; i.e., $p \vdash p$ may be weakened with anything. Further, $[\![p]\!]_P$ (the premissory role of p) is a *proper* subset of this. The things that we must adjoin to $\langle\{p\},\emptyset\rangle$ is a proper subset of those we must adjoin to $\langle\{p\},\{p\}\rangle$. What is distinctive of paradoxical sentences is their multi-facetedness. That, for example, $[\![p]\!]_P \subseteq [\![p]\!]_C$.

Now, as it turns out we can understand non-transitive and non-reflexive substructural solutions to paradox as precluding the possibility of inferential productivity for paradoxical sentences.[52] The non-reflexive solution denies the equality above: i.e. $\langle\{p\},\{p\}\rangle^{\curlyvee} = \emptyset$ at least for paradoxical sentences. This rules out the possibility of inferential productivity since nothing can be a property subset of the empty-set. The non-transitive solution denies the strict-subset relation above (i.e. $[\![p]\!]_P = \mathbf{P}$ at least for paradoxical sentences—consider that in the non-transitive setting sentences like the liar are such that $\lambda, \Gamma \vdash \Delta$ for *any* $\langle\Gamma,\Delta\rangle$). Hence, in the non-transitive setting:

$$[\![p]\!]_P = \langle\{p\},\{p\}\rangle^{\curlyvee} = \mathbf{P}.$$

I don't necessarily want to claim that either of these is a bad rendering of the situation. After all, strange sentences behave strangely. So even though it is generally true that a sentence's premissory and conclussory roles are productive (typically maximally so) with respect to one another, perhaps the strangeness of paradoxical sentences gives us reason to be suspicious of this. What I *do* want to claim is that their strangeness is not owed to this fact

[52]For non-transitive I have in mind e.g. (Ripley, 2013); for non-reflexive I have in mind e.g. (French, 2016).

about their productivity. Rather, these facts about their productivity (however they shake out) are explained by their strangeness. And it is *this claim* that I think distinguishes what I am saying from other substructural approaches to paradox. For on the account I've been sketching here, the strangeness is owed to features concerning the facets of these sentences. These are features which obtain *purely in terms of the sentence's definition* and don't rely on other features of the consequence relation. Downstream of this, i.e. whether we allow the facet to be productive and deny contraction (or rule out this possibility), lies a further substantive question. I don't want to rule out the possibility that when the cards fall, the best solution to paradox which saves a naive theory of truth might be e.g. non-reflexive or non-transitive. But I do want to emphasize the explanatory virtue of the account I've advanced.

3.2.4 Pulling the Insights Together

I've argued that many logicians (both those who endorse non-contractive logics and those who deem such logics non-sensical) are committed to the following sort of argument:

P1: (No Equivocation) In a proper accounting of reasoning, sentences must be fully disambiguated. Two occurrences of the same sentence must express the same content in the implications in which they are involved on pain of equivocation.

C: (Contraction) In a proper accounting of reasoning, implications obey contraction: if an implication with two occurrence of A as a premise is good, then the same implication with only one occurrence is good as well.

So if we are commited to No Equivocation, then we must accept contraction. If we reject contraction, we must also give up No Equivocation. This is fine, but we are only tenuously still talking about consequence at that point.

I've shown, however, that there's a gap in this argument. In particular, the following premise is needed to connect **P1** and **C**:

P2: (Sameness is Redundant) In a proper accounting of reasoning, if two sentences make the same contribution to the goodness of the implications in which they are involved, then the presence of either is inferentially redundant.

This means that we may reject contraction while maintaining No Equivocation if we reject **P2** (Sameness is Redundant) instead. I buttressed this possibility with my own account of meaning understood in terms of contribution to good implication. Such an understanding opens up logical space for the possibility of rejecting contraction while maintaining No Equivocation: that is, that two occurrences of the same sentence are fully equivalent, but they might nonetheless fail to contract: their sameness needn't be redundant.

I explained this possibility in terms of two features: multi-facetedness and inferential productivity. A facet of a sentence was understood as a kind of sub-inferential role. That is, if whenever $B, \Gamma \vdash \Delta$, we have $A, \Gamma \vdash \Delta$, then we should understand B's role as premise as a facet of A's role of premise. Facets needn't be identified with any actual sentence, but it's useful to at least introduce the idea in this way. Inferential productivity is the idea that the adjunction of one role with another produces good implications that the original sentence did not. In classical settings, for example, a sentence's premissory and conclussory roles are maximally productive with respect to one another (i.e. the schema $\Gamma, p \vdash p, \Delta$ holds).

Together these two ideas helped illuminate why a sentence might not contract. The same explanation holds for both material cases of contraction failure as well as for the paradoxes. In the latter case: I demonstrated how multi-facetedness has explanatory virtue in explaining paradoxical behavior. For these reasons we should allow violations of contraction. For these reasons we should allow violations of contraction.

3.3 Transitivity and Reflexivity

Violations of Montonicity and Contraction are two examples of ways in which, on a traditional picture of content, a sentence might appear to have divergent meanings or be ambiguous somehow. In the former case, one and the same sentence seems to imply something in one context, but not another (or be so implied); in the case of contraction, two instances of a sentence may have consequences that one instance does not (but what is the additional contribution of the second instance?). What I've argued, against this appearance is that one doesn't need monotonicity nor contraction to secure univocality of content (in one or several contexts); that is, that it is a mistake to assume that these features secure *that feature* of meaning.

Transitivity and Reflexivity are two instances of the same sort of phenomenon. The difference (from the previous cases) is that we are, in addition, considering the meaning of the sentence as it figures as a premise and a conclusion. On a traditional picture of content violations of transitivity and reflexivity seem to require some kind of ambiguity of content as the sentence appears as a premise/conclusion of an implication (in one-and-the-same implication, in the case of reflexivity; and in separate implications, in the case of transitivity). Table 4 depicts how to understand the relationship between these four structural features. That is, it is assumed that a picture of content where sentences are

Table 4: Relationship Between Four Structural Rules

	Same-Side of $\vdash$	Different-sides of $\vdash$
In the same context	Contraction	Reflexivity
In different contexts	Monotonicity	Transitivity

assigned unique contents requires that the same sentence must make the same contribution to implication regardless of where it appears (i.e. in different implications or on different sides of the turnstile). In cases where this is violated we either need divergent understandings of implication or divergent understandings of content.

I have already argued how violations of monotonicity and contraction don't require such divergence if we understand content in terms of contribution to good implication. Instead, we are able to understand that one and the same content may be involved in an implication relation which is non-monotonic and non-contractive without appeal to a deeper layer of content to explain this substructural behavior. I argued that understanding content in this way means pushing the substructurality of implication all the way down into the content. By contrast, the postulation of a further level of content (which behaves fully structurally), I labelled the assumption of structurality.

What I argue in this section is that transitivity and reflexivity are also not required to secure this feature of meaning. That is, if we understand the content of a sentence in terms of that sentence's contribution to good implication, then we may have a sentence which is both non-transitive and non-reflexive and this behavior can be understood without appeal to some further layer of content. To argue for this I look at an example of a non-transitive logic and an example of a non-reflexive logic. I try to explain the way in which this substructural behavior is explained. In particular, I argue that in both cases the substructurality is not understood to be a feature of the content itself.

3.3.0.1 Non-Transitivity

First, let's look at an example of a non-transitive logic. In particular, I'll discuss so-called "Strict-Tolerant (ST)" Logics.[53] The idea is fairly simple. Take strong Kleene logic and define a notion of strict-tolerant entailment as holding whenever all premises are "strictly true" and no conclusion is "strictly false". In classical semantics (i.e. two-valued semantics) this is homophonic with the traditional notion of entailment. But with a third truth value, it can diverge from what one might think the traditional notion was (usually: the conclusion is true if all the premises are true).

Definition 3.3.1 (ST-Model)**.** A model M consists of a valuation function $v(\cdot)$ which maps sentences of the language $\mathcal{L}$ to truth values, $v : \mathcal{L} \mapsto \{1, \frac{1}{2}, 0\}$ (1 is true, $\frac{1}{2}$ is neither-strictly-true-nor-strictly-false, 0 is false).[54]

Definition 3.3.2 (ST-Entailment)**.** We say that α *strict-tolerantly* entails β in M, written $\alpha \vDash_M \beta$ iff: if $v(\alpha) = 1$, then $v(\beta) \gneq 0$.

We say that α entails β, written $\alpha \vDash \beta$ iff whenever a valuation has it that $v(\alpha) = 1$, we have $v(\beta) \gneq 0$, (i.e. $\alpha \vDash \beta$ for all M).

Now suppose $\alpha \vDash \beta$, and $\beta \vDash \gamma$. How can it happen that $\alpha \nvDash \gamma$? We simply need to find a counter-model for this last entailment that is not a counter-model for the other two. But this is simple. Suppose we have M such that $v(\alpha) = 1, v(\beta) = \frac{1}{2}$, and $v(\gamma) = 0$. Clearly

[53]ST was originally developed by van Rooij to deal with vagueness. It has since been used for a number of purposes. See Cobreros et al. (2012a,b); Barrio et al. (2015a).

[54]How to understand the third truth-value in ST is not uncontroversial—or so it seems to me. Some have argued that ST is extremely similar to Graham Priest's Logic of Paradox (LP)—introduced in Priest (1979)—in which case $\frac{1}{2}$ is a glut (both true and false). But Ripley (2013), for example, seems to suggest a different understanding of this third value.

$\alpha \vDash \beta$ and $\beta \vDash \gamma$. Equally clearly: $\alpha \nvDash \gamma$.

How should we understand this phenomenon? As I remarked above, transitivity can be understand as the insistence that a sentence *means the same thing* as a premise and as a conclusion. And here that is precisely violated.

$$\begin{aligned}
\alpha &= \text{“Chris didn't eat cake”}, \\
\beta &= \text{“Chris didn't eat too much cake”}, \\
\gamma &= \text{“Chris ate cake”} = \neg\alpha.
\end{aligned}$$

I also suppose that there is the relevant logical connection between β and the other sentences (I leave the details ambiguous, but this is to ensure that the valuation function respects the relation in meaning between the sentences). Clearly the valuation above seems to get many details right including the fact that β might be neither true nor false. What I want to argue for is that this picture understands β to be saying something different as a premise and conclusion.

Now, what, in general $\alpha \vDash \beta$ *means* is that if α is assertible, then β is “tolerably-assertible”. This is to say—literally—that an assertion of β is tolerable; asserting β is not ruled out by having asserted α. So when we consider $\alpha \vDash \beta$, we must consider β as “not ruled out”. This is intuitively what the semantics does. The entailment holds just in case $v(\beta) \in \{1, \frac{1}{2}\}$ when $v(\alpha) = 1$.

But what happens when we consider $\beta \vDash \gamma$. Here we have it that if β is assertible then γ is tolerably assertible (again: in the sense that it is not ruled out). So, again, the reason that $\alpha \nvDash \gamma$ is precisely that while α makes β tolerably assertible, β's tolerable assertibility says nothing about γ's; only β's assertibility does this. In fact, it is possible that β's tolerable assertability does not even guarantee γ's tolerable assertability, since there is a valuation where $v(\beta) = \frac{1}{2}$ and $v(\gamma) = 0$.

When we want to determine whether $\alpha \vDash \beta$, we focus our attention on a range of models (namely those where $v(\alpha) = 1$. Since we supposed this entailment holds, we know that $v(\beta) \gnsim 0$. The traditional account would be happy with this result. It's certainly not false that “Chris didn't eat too much cake”, but we don't want to say that it's “strictly true” either, at least not generally.

However, when we want to determine whether $\beta \vDash \gamma$, we focus our attention on a range of models (namely those where $v(\beta) = 1$, that is, we ask when β is strictly true (assertible, full stop), is γ tolerably true?

Notice that what β means, in the sense of which models we consider (in which situations we can assert it) varies when we consider β as a premise and β as a conclusion. β as a premise means something different than β as a conclusion, and this is precisely what a failure of transitivity amounts to.

What I want to suggest, or what I want to make room for, is the thought that β does not mean something different. When someone says "Chris didn't eat too much cake" (tolerably so following the assertion that "Chris didn't eat cake"; full-stop preceding the assertion that "Chris ate cake") they are not saying two different things. ST, which makes room for a non-transitive consequence relation, nevertheless maintains much of the assumptions of traditional semantics. It therefore must interpret sentences divergently in order to get the proper inferential behavior.

3.3.0.2 Non-Reflexivity

Finally, I'll examine an account of non-reflexive consequence put forward by French (2016). The notion of consequence here is based on "q-consequence" (=quasi-consequence) from Malinowski (2004, 2014). French writes:

> "Let us read a sequent $\Gamma \succ \Delta$ as telling us that if we *do not reject* all the members of Γ then we should *accept* some member of Δ. What does this reading of sequents tell us about the meaning of [*Id*] and [*Cut*]. On this account [*Id*] tells us that if you don't reject C then you should accept it. So reflexivity will fail for any sentence C which one should neither reject nor accept." (French, 2016, p. 122)

French argues that we should understand paradoxical sentences as sentences on which we cannot take a stand. The reason that we can't take a stand, the thought goes, is that these sentences are undetermined in some significant way, which is to say that they don't express determinate thoughts. What is wrong with paradoxical reasoning, French argues, isn't any particular step in the chain of reasoning, but rather the initial assumption that we can take a stand on all sentences. Notice that when Γ and Δ are singletons (say $\Gamma = \Delta = \{p\}$), then satisfaction of reflexivity is simply bivalence: if it's not rejected, it must be accepted. Equivalently: it's not possible to not-reject all premises and not-accept all conclusions, hence $p \vdash p$ says: it's not possible to both not-reject and not-accept p.

We can also provide a 3-valued semantics for this non-reflexive logic. This logic is often called "TS" for tolerant-strict logic. This is because implication involves tolerant truth of the premises (i.e. not-false) and strict-truth of the conclusions. As with ST above, this is homophonic with classical logic in a two-valued setting, but given a third value this has the potential to come apart, here are the formal details.

Definition 3.3.3 (TS-Model)**.** A model M consists of a valuation function $v(\cdot)$ which maps sentences of the language $\mathcal{L}$ to truth values, $v : \mathcal{L} \mapsto \{1, \frac{1}{2}, 0\}$ (1 is true, $\frac{1}{2}$ is neither-true-nor-false, 0 is false).[55]

Definition 3.3.4 (TS-Entailment)**.** We say that α *tolerantly-strict* entails β in M, written $\alpha \vDash_M \beta$ iff: if $v(\alpha) \gneq 0$, then $v(\beta) = 1$. Or: it is not the case that $v(\alpha) \neq 0$ and $v(\beta) \neq 1$.

We say that α entails β, written $\alpha \vDash \beta$ iff whenever a valuation has it that $v(\alpha) \neq 0$, we have $v(\beta) = 1$, (i.e. $\alpha \vDash \beta$ for all M). Or, equivalently: there is no model with $v(\alpha) \neq 0$ and $v(\beta) \neq 1$.

Before proceeding, I'd like to prove one small result for TS, namely that in TS, cut is a valid sequent rule. I make use of this result in the section following this.

Proposition 3.3.5. Transitivity is a valid meta-rule in TS. That is:

$$A, \Gamma \vdash \Delta; \Gamma \vdash \Delta, A \Rightarrow \Gamma \vdash \Delta.$$

Proof. For expediency, I prove a much simpler case, namely: $A \vdash B; B \vdash C \Rightarrow A \vdash C$. I prove this semantically. Suppose that there are no valuations with $v(A) \neq 0$ and $v(B) \neq 1$, and likewise none with $v(B) \neq 0$ and $v(C) \neq 1$. We want to show that there can be no valuation with $v(A) \neq 0$ and $v(C) \neq 1$. By way of contradiction suppose there was, i.e. suppose $v(A) \neq 0$ and $v(C) \neq 1$. Then by supposition such a valuation would require $v(B) = 1$ (since there is no valuation with both $v(A) \neq 0$ and $v(B) \neq 1$) and would also require $v(B) = 0$ (since there is no valuation with both $v(C) \neq 1$ and $v(B) \neq 0$). But then

[55] It seems right to say that $\frac{1}{2}$ is a gap.

$v(B) = 1$ and $v(B) = 0$, a contradiction. Hence there can be no valuation with $v(A) \neq 0$ and $v(C) \neq 1$, thus $A \vdash C$. ■

But given these definitions, we cannot generate any consequence relation.[56] Supposing p is atomic, for example, there will be a valuation where $v(p) = \frac{1}{2}$ and hence a model where $v(p) \neq 0$ and $v(p) \neq 1$, hence $p \nvDash p$. In order for a sequent to hold, it must be the case that some element never receives value $\frac{1}{2}$, but there is no way to rule this out.

The reason that we don't get $p \vdash p$ (for any sentence) is that there is no guarantee at this level of abstraction that p has a determinate content. The best we can express, for French is what would follow *if we assumed that p did* (have such a content). That is, we only get the validity of meta-sequents.

> "Given this understanding of how we should read individual sequents we are now faced with explaining metasequents. Mostly, in this setting, we are interested in how we should understand metasequents which have the following form:
>
> $$P_1 \succ P_1, \ldots, P_n \succ P_n \Rightarrow \Gamma \succ \Delta$$
>
> My proposal is that we read metasequents of this form as telling us that if we take a stance on each of the P_is, then if we don't reject all the members of Γ then we should accept some member of Δ, where to 'take a stance on A' is to either accept A or to reject A.
> The intuition behind this reading is to think of metasequents as making explicit the fact that in judging a sequent valid we are **taking its components to be the kinds of statements for which it is permissible to take certain attitudes towards**. This allows us to recapture full classical reasoning as being enthymematic, involving a suppressed assumption that the statements involved are ones which we either accept or reject." (French, 2016, p. 123, bold mine)

[56]This relies on using "strong Kleene" truth tables for the connectives (which French does):

$\to$	0	$\frac{1}{2}$	1
0	1	1	1
$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	1
1	0	$\frac{1}{2}$	1

	$\neg$
0	1
$\frac{1}{2}$	$\frac{1}{2}$
1	0

Different connective definitions would allow for some consequences to hold. For example Łukasiewicz 3-valued connectives would still yield violations of reflexivity for French (namely of atomic sentences), but you would get some true sentences and thus true consequences (which would also bring French's system closer to other supervaluationist approaches, if we can call his approach such):

$\to$	0	$\frac{1}{2}$	1
0	1	1	1
$\frac{1}{2}$	$\frac{1}{2}$	1	1
1	0	$\frac{1}{2}$	1

	$\neg$
0	1
$\frac{1}{2}$	$\frac{1}{2}$
1	0

Notice that $v(p \to p) = 1$ on all valuations. This is perhaps unsurprising since part of the motivation of Łukasiewicz is precisely to allow the indeterminate to imply the indeterminate.

If we assume that sentences have determinate contents, then certain further things may follow (but we need not assume that).[57]

I think we can see this sort of thought more clearly by considering a more recent account of non-reflexive consequence (also given in response to paradox). Nicolai and Rossi (2018); Murzi and Rossi (2021) construct a hierarchy of *sets of sequents* in order to reach a Kripke fixed point. This fixed point yields two desired results: first it gives us a consequence relation; second it is a consequence relation which has a naive validity predicate (by which is meant, roughly, that $V(\ulcorner\alpha\urcorner, \ulcorner\beta\urcorner)$ holds iff $\alpha \vdash \beta$. Of course, because we are interesting in implication relations, we should formulate this in terms of one. In particular, they[58] prove that this hierarchy gets us:

$$\Gamma_0 \vdash \phi, \Delta_0 \text{ and } \Gamma_1 \vdash V(\ulcorner\phi\urcorner, \ulcorner\psi\urcorner), \Delta_1 \Leftrightarrow \Gamma_0, \Gamma_1 \vdash \psi, \Delta_0, \Delta_1, \tag{VDm}$$

which is a kind of validity discharge rule. Likewise:

$$\Gamma, \phi \vdash \psi, \Delta \Leftrightarrow \Gamma \vdash V(\ulcorner\phi\urcorner, \ulcorner\psi\urcorner), \Delta. \tag{Val}$$

However, this consequence relation generated has some violations of reflexivity. All other structural rules hold unrestrictedly. Further they are able to say when reflexivity is violated. If $p \nvdash p$ it is because p is not "grounded". Groundedness is defined as follows:

> "For every ordinal α and every sequent $\Gamma \vdash \Delta$, if $\Gamma \vdash \Delta$ is in I^α_ψ [the hierarchy of sequents at level α], then there is at least one setnence ϕ in Γ such that $\phi \vdash \emptyset$ is in I^α_ψ, or at least one sentence in Δ such that $\emptyset \vdash \psi$ is in I^α_ψ." (Nicolai and Rossi, 2018, p. 568)

[57]In rehearsing Humberstone's account of possible non-reflexive consequence (Humberstone, 1988), French offers the following example:

> "Consider, then, a case where both premises and conclusion are taken from the same language, say English, but are assessed relative to different dialects: for concreteness let us have the premises assessed as true relative to Australian English and the conclusion's relative to British English. Then the argument from 'John lost a thong at the beach' to 'John lost a thong at the beach' will count as a counterexample to structural reflexivity when John lost a flip-flop at the beach, but didn't lose any swimwear." (French, 2016, p. 119)

[58]Rossi is a co-author on both pieces. There are some minor differences in the aims of the papers, but the technical details are nearly identical. At the level of detail I am rehearsing, differences can be ignored.

To appreciate this, consider that the first step in the hierarchy are sequents of the form:

$$\Gamma, \phi \vdash \Delta,$$
$$\Gamma \vdash \psi, \Delta$$

where ϕ is some concrete falsehood and ψ some concrete truth.[59] The result of all of this is that we get violations of reflexivity only in particular cases: where the sentence is not grounded. To be not grounded just means that the sentence is not given a determinate truth-value, i.e. its truth is un(der)determined (for French this meant that we needn't take any stance on it).

In a properly regimented language, sentences always receive contents which are determinate enough to receive truth values and thus on which we can take some stand. It is only defective sentences, which fail to be determinate in this way. For example, the liar $\lambda =_{df.} \neg T(\lambda)$ or the validity-Curry[60] $\pi =_{df.} V(\top\pi, \ulcorner\bot\urcorner)$. Both of these sentences are *ungrounded* because we can't trace them back to some determinate truth or falsehood. Both the derivation of the liar and the derivation of Curry start with a sequent of the form:

$$p \vdash p$$

which is an instance of reflexivity. But as we were constructing the hierarchy, we only get a sequent of this form if one of the following holds:

$$\emptyset \vdash p,$$
$$p \vdash \emptyset.$$

which is to say that p has some determinate truth-value.

The thought then is that because some sentences do not express a single, determinate thought, such sentences do not receive truth values and thus cannot even imply themselves (if implication is taken as a constraint on what attitudes we can or cannot hold with regard to a *sentence*). If we precisify such a sentence (so that it expresses a single, determinate thought),

[59]Nicolai and Rossi (2018) use Peano Arithmetic to generate truths and falsehoods of the form $s = t$ (where $s, t \in \mathbb{N}$). Murzi and Rossi (2021) start with a valuational model $\mathcal{M}$ and use $\mathcal{M} \models \psi$ and $\mathcal{M} \not\models \phi$ to generate initial sequents.

[60]Or replace '$\bot$' with some arbitrary falsehood.

then the precisification (the thought that the sentence *would* express, so precisified) could be used in reasoning. This is what French means when he invokes metasequents. A metasequent tells us logical relationships between sentences should they be precisified appropriately.

Parenthetically: in the next chapter, I explore a view similar to this one in the philosophy of language put forward by Charles Travis.[61] Travis' view is that sentences do not express determinate thoughts except when they are thought on particular occasions. We can think of sentences as offering something akin to a description, but a description isn't a commitment regarding a concrete way things could be. A description is (perhaps) a kind of rough procedure for how to reach a thought. Thus one and the same sentence can be said truly on one occasion and said falsely on another, because the thought expressed on one occasion differs from that expressed on another. This view will itself be precisified in the next chapter. I mention it now to help give shape to the view I am describing now.

Tightly associated with this sort of view is an acknowledgment that sentences may stand in substructural relations of consequence. They do so because a sentence on its own does not yet express a determinate thought. When a sentence is precisified or thought on a particular occasion, it does express a thought and *relative to that occasion/precification* behaves perfectly structurally. This is precisely what I have been calling the assumption of structurality throughout this chapter. My suggestion, by contrast, has been to understand a view of content where violations of structural rules inhere at the "deepest" level. That is, that we push the substructurality all the way down into the content, i.e. into the precisification. The assumption that thought, properly precisified cannot do so, is precisely what I want to contest.

3.3.0.3 Some Lessons

In this section I exhibited ways in which non-transitive and non-reflexive logics presuppose the assumption of structurality. In many ways non-transitive and non-reflexive logics do not differ drastically from the logics surveyed earlier. That is, the way in which the assumption of structurality manifested should be familiar. A violation of a structural rule (on

[61](Travis, 2000, 2008)

this assumption) seemed to require the possibility of further determination of the content of a sentence which could either vary within a single context (i.e. in a single sequent: and this led to the violation of contraction and reflexivity) or at least from context to context (i.e. between different sequents: and this led to the violation of monotonicity and transitivity).[62]

I'd like to close this section by observing the ways in which the non-transitive and non-reflexive approaches surveyed are deeply related.[63] In both cases, what seems to underlie the idea that consequence can be non-transitive and non-monotonic is the thought that while sentences may stand in these relations, the content that those sentences express do not. In particular, when we precisify the sentences properly, non-transitive and non-reflexive consequence simply falls away. The difference between ST and TS concerns what we allow ourselves to do with unprecisified sentences.

If we're open to *tolerating* such sentences (providing some precisification exists), then we will allow ourselves to reason with the unprecisified sentences. Notice what $p \vdash_{ST} p$ *says*: that it is not possible that $v(p) = 1$ and $v(p) = 0$, that is either $v(p) \neq 1$ or $v(p) \neq 0$: some precisifcation exists, p is either not strictly true or not strictly false (or: p is either tolerably true or tolerably false, but we needn't be more precise than that). But that we can't be more precise than that precisely rules out a guarantee that p is precisified the same way across contexts. ST is tolerant of such sentences. By contrast, TS says we cannot reason with a sentence until it has been precisified. $p \vdash_{TS} p$ *says* that it's not possible that $v(p) \neq 0$ and $v(p) \neq 1$, that is either $v(p) = 1$ or $v(p) = 0$ (or: p is either strictly true or strictly false, anything less won't be *tolerated*). Unless and until we've precisified a sentence (isolated the exact content we wish to express), we cannot reason with it at all, and thus we do not even get reflexive consequence.

Both of these logics, of course, are introduced to solve the same sort of problem: para-

[62] A recent paper by Cobreros et al. (2017) observes something similar; i.e. a deep similarity between non-monotonicity and non-transitivity.

[63] To those familiar with the substructural literature on paradox, I think such familiarity is not terribly surprising (if only for sociological reasons). I think it can be best appreciated by considering the analogy:

ST : subvaluationism :: TS : supervaluationism

I don't want to claim that ST is subvaluationist. Though many have claimed that that is precisely what it is, i.e. that it's simply LP—this debate has generated its own literature, but see Barrio et al. (2015b, 2016) for the initial argument (in response to Cobreros et al. (2012b)); see Cobreros et al. (2020) for a recent reply.

doxes of transparent truth.[64] In this case, they are interested in whether paradoxical sentences like the liar can be precisified: whether we can think a single, definite content with them. In both cases, no such content is to be found, hence either we can never rely upon it between contexts of reasoning (from one sequent to another: as ST-logicians have it), or we can never reason with it at all (as TS-logicians have it).

What I am advocating is the idea that we push such behavior down into the content itself. It is possible for a sentence to express a single, determinate content and still stand in such consequence relations. This means that a sentence need not have determinate truth conditions to express a content. Even vague or paradoxical sentences may be understood in this way. The fact that we can reason with a sentence, in a variety of contexts, ought to be proof enough that the thing we consider when so reasoning is contentful. What I am suggesting is taking our ability to reason with such sentences seriously: that is, to say that we don't equivocate or make some sort of mistake, but that it is just a feature of the content of vague or paradoxical sentences that they are substructural. In the case of a paradoxical sentence like the liar, this feature is *formal*, in the sense that we can characterize it formally using methods of logic. Vague sentences (and similar such things) have these features materially.

3.4 Conclusions

This chapter articulated the ways in which my approach to the semantics of sentences that stand in substructural relations of implication differs from other approaches in the literature. In particular, I focused on the way in which my approach involves *pushing the substructurality all the way down in the content*, by understanding substructural implication to be a part of the "deepest" level of a sentence's meaning. Approaches which failed to do this—usually via the postulation of a further level of content which behaves fully structurally—were characterized as implicitly accepting what I called the *assumption of*

[64]ST is also leveraged to deal with vagueness and related phenomena. I'm unaware of similar work on TS, though there are of course supervaluationist solutions to vagueness.

structurality. The assumption of structurality typically took the following form. Suppose a sentence A is involved in a pattern of substructural implication, which I'll characterize as:

$$\begin{aligned} A, \Gamma_i &\vdash \Theta_i \\ A, \Lambda_i &\nvdash \Delta_i \\ \Xi_i &\vdash \Pi_i \\ \Phi_i &\nvdash \Psi_i \end{aligned}$$

for $1 \leq i \leq n$ (there's some combination of sequents where A does/does not occur and which do/don't hold which characterizes substructural behavior). Then while the sentence A does stand in these relations, this is explained by a further level of content which does not. That is, there are a number of contents that A may express or through which A may be made more precise such that:

$$\begin{aligned} A_i, \Gamma_i &\vdash \Theta_i \\ A_i, \Lambda_i &\nvdash \Delta_i \\ \Xi_i &\vdash \Pi_i \\ \Phi_i &\nvdash \Psi_i \end{aligned}$$

But since each of these A_i are distinct, there is no *actual* substructural implication (at least at the deepest level of content), since there is no A_i that occurs in all of the sequents that generated the substructural pattern of implication to begin with. The difference between the ways that this postulation of a further level content arise were explored throughout the chapter. A general lesson concerning their relationship was illustrated in Table 4 (reprinted here as Table 5). The semantics of substructural content that I provided, by contrast, did not involve an acceptance of the assumption of structurality. I insisted that we push the substructurality all the way down into the content.

Along the way, there were three separable contributions that I made worth highlighting. First, my understanding of non-monotonic content paves the way for understanding how my account can help to open up logical space in a number of different debates in philosophy. In the next chapter, I look at this in more detail.

Table 5: Relationship Between Four Structural Rules (Redux)

	Same-Side of $\vdash$	Different-sides of $\vdash$
In the same context	Contraction	Reflexivity
In different contexts	Monotonicity	Transitivity

Second, my examination of non-contractive content located a central joint in the literature on substructural solutions to paradox. That joint is the idea that a sentence may be multi-faceted. As I argued, we can understand what unites the different substructural solutions to paradox (in particular, non-contractive, non-transitive, and non-reflexive) in terms of the idea that content is multi-faceted. Then what distinguishes those solutions from one another is whether (and if not, why not) they think this multi-facetedness is inferentially productive. I argued that what is distinctive of paradoxical sentences (and does the bulk of the explanatory work) is something that all of these accounts should be able to agree upon: the multi-facetedness of the sentence.

Finally—continuing the theme of understanding the way that different substructural logics may be connected—is a better appreciation of the connection between non-transitive and non-reflexive solutions to paradox. When exploring non-contractive solutions, I argued that these two solutions belong in a box as denying that paradoxial sentences have *inferentially productive* facets (so that the central disagreement between non-contractive, on the one side and non-transitive and non-reflexive, on the other, concerns inferential productivity). In the last section, I argued that we can understand the way that non-transitive and non-reflexive logics are related to each other.[65] That is: when dealing with sentences that allow multiple disambiguations/precisifications, must we first precisify the sentence to reason with it, or can be reason with the sentence unprecisified? Non-reflexive logics (TS) answer the former question affirmatively (full precisification is required); non-transitive logics (ST) answer the latter affirmatively (no such precisification is required).

[65]Parenthetically, I suggested this looks a lot like the relationship between subvaluationism and supervaluationism.

4.0 The Significance of a Consideration

So far I have argued that the assumption of structurality is not required to furnish an adequate notion of content. If we understand the content of e.g. a sentence in terms of the contribution it makes to good implication, a fully adequate notion of content emerges in which sentences may stand in relations of consequence that need not be fully structural (i.e. needn't be monotonic, transitive, reflexive, or contractive).

The goal of the present chapter is to show that this notion of content can do more than satisfy philosophers who were already sympathetic to inferentialism or substructural logic. That is, that the notion of substructural content can be used to do real work in a broad range of philosophical debates. I do this by showing that it can help clear up two debates in disparate regions of philosophy: radical contextualism in the philosophy of language and particularism in meta-ethics.

The radical contextualism and semantic minimalism debate in the philosophy of language concerns whether there is such a thing as semantic content, i.e. whether there is something said literally by a sentence when asserted (ignoring the influence of extra-sentential or contextual information). As I'll explain below, semantic minimalists hold that sentences express minimal propositions (i.e. literal meanings which do not vary with the context of assertion). Contextualists hold that what is said by a sentence on a particular occasion may be influenced (perhaps radically so) by contextual information. A practical consequence of this debate is this: if I quote someone out of context, do I risk misattributing a claim to her (i.e. saying of her that she said something which she did not) or do I only risk conveying something false (namely that she might have meant something that she did not mean)?

The moral particularism and moral generalism debate in meta-ethics concerns whether normative verdicts—i.e. verdicts about whether an act ought or ought not be performed—are governed by moral principles. A moral principle is a connection between various considerations and an action (or a normative verdict concerning it). Generalism, in its purest form, holds that it is features intrinsic to an act that determine whether that act ought to be performed (i.e. is mandatory/prohibited/permitted). More moderate forms may allow that features extrinsic to an act can have some effect on normative verdict. Particularism, in its most extreme form, denies that there are any moral principles whatsoever. That we can only speak of normative verdicts *on particular occasions* and that the way that various features of an act and extrinsic considerations bear on normative verdicts cannot be stated with any generality.

In addition to making (what I take to be) a substantive contribution to each of these debates, I intend to make a more general and potentially separable contribution connecting these topics by demonstrating a deep, and thus highly illuminating, commonality between them. It is easy for me to say now that this insight strikes me as incredibly obvious, and there are ways of stating it that make it seem so: a) that radical contextualism is just particularism viz. truth conditions; or b) that particularism is just radical contextualism concerning normative verdicts or deontic judgments. Anyone with a superficial understanding of these debates can probably get an intuition about how that might shake out, but, as with many things: while the obvious statement isn't wrong, what it lacks in nuance it also lacks in insight. Part of the upshot here is finding that nuance and providing that insight. This separable contribution highlights how things hang together.[1]

That is to say, that while (a) and (b) aren't wrong, they aren't great statements of the commonality. While I won't be able to fully illuminate the connection until the end of the chapter, I'd like to prime the reader by providing a rough guide of how things in each debate line up. [To the reader: feel free to return to Table 6 to check your intuitions. If you are unfamiliar with these debates, I suspect this chart may not be terribly helpful yet.] I hope to explain how the items in the left-hand column correspond to those in the right-hand column by the end of this chapter. Here is a further preliminary statement to help guide my reader:

P1: There can be such a thing as the content of a sentence (or a moral principle) only if there are jointly sufficient and individually necessary conditions concerning the significance of a consideration to e.g. the truth conditions of an assertion, or deliberation concerning

[1] I'm not aware of anyone in the contextualism/minimalism debate discussing particularism, but Dancy (2004b) basically acknowledges (a), when he acknowledges that the rules for determining the "meaning" of a sentence (or expression) are sensitive to the same sort of considerations (pp. 194ff.).

Table 6: The Components of the Two Debates

Radical Contextualism & Semantic Minimalism	Moral Particularism & Moral Generalism
Sentence	Action
Truth/falsity of an assertion	(Overall) normative verdict
Truth-Conditions	Reason(s)
Semantic Content	Moral Principles
Contextual Information	Contributory Reason
Radical Contextualism	Moral Particularism
Semantic Minimalism	Moral Generalism

action.[2]

P2.a: There can be such a thing as the content of a sentence (or a moral principle) [Affirming the Antecedent]

P2.b: There are **no** jointly sufficient and individually necessary conditions concerning the significance of considerations. [Denying the Consequent]

My view is that all participants of both debates endorse some version of P1. Radical contextualists and particularists endorse some version of P2.b,[3] while minimalists and generalists endorse some version of P2.a. I agree with P2.b for reasons sketched previously in the dissertation, but I think that radical contextualists (and particularists) go too far in abandoning the idea of the content of a sentence or rules for action. Hence, a kind of flat-footed intuition underlying minimalism (and moderate contextualism) as well as generalism, is right: there are principles for action and there are contents of sentences. The issue is the assumption

[2]This will be made more precise for both debates and in general in the course of the paper. I'm just trying to prime the reader for what's to come.

[3]So-called "moderate contextualists" are a middle course that will require some care. On the minimalist's telling, all forms of contextualism collapse into radical contextualism. I briefly explain below the problems with this view. Instead, it seems clear that the dividing line is between radical contextualists and everyone else (thus moderate contextualists belong in a box with minimalists not with radical contextualists).

of P1: that such things require a certain form. A form that would be guaranteed if the assumption of structurality holds. I argue that that assumption does not hold in general and should not be imposed by our theories.

This chapter is divided into three parts. First (§4.1), I go through the debate in radical contextualism. Next, in (§4.2), I go through the particularist debate. My intention is that each of these sections can be read relatively independently of each other. Finally (§4.3), I draw together the insights and explain the commonality of each debate.

4.1 Radical Contexutalism and Semantic Minimalism (and also "moderate" contexutalism)

The debate[4] concerning radical contexualism and semantic minimalism turns on a very small question: whether, in general, what is said by the utterance of a sentence (i.e. the propositions expressed) may be determined solely (or: for the most part) by that sentence's meaning. Semantic minimalists think (with a few minor exceptions) that what is said by a sentence is simply a matter of that sentence's meaning.[5] Radical contextualists deny that sentence meaning on its own could play even a merely constraining role on what is said by a sentence. It will help to clarify some of these distinctions.

[4]Throughout this section I relied heavily on Bezuidenhout (2017); Borg (2007) for background. When characterizing Semantic Minimalism I had primarily in mind Borg (2004); Cappelen and Lepore (2008). When characterizing radical contextualism, Travis (2000, 2008), though Searle (1969, 1983) is also an important figure in this debate. The moderate contextualist that I've relied on most is Recanati (2004). In fact, though I disagree with much of his view, I find it nearly impossible to remove his influence on my understanding of the debate. In addition to Recanati, there is also a wide diversity of views which could be counted as moderate contextualist, especially if we include relevance theorists working in linguistics (Sperber and Wilson, 1995; Bezuidenhout, 2009).

[5]There is a giant asterisk that must be placed on this claim. Many minimalists think that speaker intention (perhaps most obviously as it figures in demonstration) plays a large role in determining exactly what proposition in expressed. For example, if I utter an ambiguous sentence, we can ask whether the context settles *what I actually said* or whether *what was said* is a matter of something in my head (perhaps the felicity of my assertion is somehow constrained by context). Of course speaker intention is, or can be, extremely opaque. So anything the radical contextualist says about content could be hoisted upon speaker intention. I think the minimalist would agree largely with this, which is why she and the radical contextualist are not so far apart after all. The minimalist just thinks there's a small island of semantic content that plays a non-eliminable role in e.g. pragmatic reasoning.

While the issue of speaker intention is surely important—and something that minimalists need to account for—I largely ignore it.

Perhaps the first distinction we must get clear on is that between "what is said" by a sentence and "what is meant" by a sentence.[6] Grice introduced this distinction[7] in order to distinguish the content of an assertion of a sentence from what we mean by asserting that sentence. For example, if I assert the sentence, "I am tired", then what is said by my assertion is that I, Dan Kaplan, am tired, but on a particular occasion, I might mean *that I need another cup of coffee*, or *that it is time for my guests to leave.* In general, *what one means* by asserting a sentence can be quite diverse. The content of this may be called the pragmatic content of the assertion; the content of the former (what is said), the semantic content (hence pragmatics aims to uncover the structure of the latter; semantics that of the former). To help keep things clear and in line with widespread convention, I'll call *what is meant, but not said*, "what is implicated". The significance of what is meant but not said (what is implicated) is that this content of an assertion can be false without affecting the veracity of the assertion. By contrast, one has *said somethings false* if what is said turns out to be false; and of course this difference is significant not just in the philosophy of language, but in many other domains as well.[8] For my purposes, the difference concerns the speakers relationship to these various contents. Whether the speaker has spoken falsely and what claims she can cancel (=clarify as not being meant) without revising what she has said.

While *what one means* by the assertion of a sentence can be quite diverse, the opposite is typically taken to be the case for *what is said* by the assertion of the sentence. Grice, for example, took the latter to be wholly conventional, and the semantic-pragmatic distinction often sees conventionality in the former (but not the latter).[9] Further, this sort of stability

[6] I am using "what is meant" to correspond roughly to "what is implicated" by a sentence. I've chosen this locution because I think the pragmatic content of an assertion can often overlap with the semantic content and I want the distinction to capture this. So I might mean what I say (but, at least on the standard Gricean picture, I can't implicate what I say).

"What is said" is meant to stand in for the "semantic content" of a sentence, *strictly speaking.* This way of cashing it out lines up most closely with how Recanati (2004) cashes things out. As I'll elaborate below, the central question is: "what determines the semantic content of a sentence when uttered on a particular occasion". I'm treating this as equivalent to "what is said" by the assertion of the sentence. I treat Semantic Minimalism as the view which says that what is said is always the literal meaning of a sentence.

[7] Grice (1957, 1975, 1991).

[8] What is implicated by a sentence maintains some degree of plausible deniability, for example. In a legal context, this can make the difference concerning whether an assertion constitutes defamation.

[9] Of course, Grice notes that there are "conventional implicatures", i.e. things which seem to be implied by the assertion of a sentence (and so could be meant by its assertion) which are a matter of convention. Whether there are such things is controversial and not of interest to the aim of the present chapter.

concerning *what is said* is often taken to be important for determining *what is meant* by a sentence, since determining what is meant often requires knowing *what is said*. For example, when I say "I am tired," having just looked down at my empty cup of coffee, it is reasonable to conclude that what I mean is that I need another cup of coffee; having just looked at my watch, that it is time for my guests to leave; and having been asked how I'm feeling, I might just mean that I am tired. But in all these cases, what I say—that I am tired—seems to play a crucial role in determining what I mean.

Nevertheless, it often seems to be the case that what is said by the assertion of a sentence cannot be settled wholly by convention or by the sentence's meaning. For example, I might say "I've had breakfast" not to say that *at some point* in my past, I've had the meal, but rather to say that *today*, I've already eaten breakfast. If in response to a friend proposing that we get some breakfast, I reply, "I've had breakfast", my meaning is clear: I am turning down their proposal. But if I had not eaten breakfast on that day, is what I said false, or just deceptive? And of course, it is possible to say that one has had breakfast at some point in the past, which is why it seems that what is said by the assertion is not settled strictly by the sentence's conventional meaning. For example, many people skip breakfast entirely and it might be reasonable for them to say "I've had breakfast," when trying to explain their dietary choices.[10] It is easy to generate sentences like this.[11]

To help keep the different positions in this debate distinct, it will be helpful to introduce another term of art, *literal meaning*.[12] The literal meaning of an utterance is *what is said* by an utterance only with respect to that sentence's (conventional) meaning. The literal meaning of a sentence is what would be said by that sentence if no (or *as little as possi-*

[10]–What do you mean you skip breakfast? Do you mean to tell me that you've never eaten pancakes? Waffles? Scrambled eggs? –Firstly, it is possible to have those foods outside of breakfast, but also: I've had breakfast.

[11]A list from Recanati (2004, p. 8) includes:

1. I've had breakfast. [today?]
2. You are not going to die. [ever?]
3. It's raining. [here?]
4. The table is covered with books. [there exists a unique table such that...?]
5. Everybody went to Paris. [everyone?]
6. John has three children. [only or at least?]

[12]Outside of this debate the contrast of literal meaning is metaphorical meaning. Here literal meaning is contrasted with various other non-figurative things that a sentence might mean.

ble—perhaps, "minimal") determination outside of the sentence's meaning is required. For example, the literal meanings of the sentences in Table 7 can be gleaned from their paraphrases.[13] The literal meanings of sentences are often not transparent to speakers who have

Table 7: Sentences and their Literal Meanings (paraphrased)

Sentence	**Literal Meaning**
"I've had breakfast"	At some point in the past, I've eaten food before noon
"You are not going to die"	You will live forever
"It's raining"	Liquid precipitation is occurring somewhere
"The table is covered with books"	There is a unique x such that x is a table and covered with books
"Everybody went to Paris"	Every person, living or dead, has gone to Paris at some point before now
"John has three children"	The number of children John has is greater than or equal to three

not studied linguistics or the philosophy of language. In fact, it often takes some practice to be able to find them. Further, *what is said* by a sentence on a particular instance often seems to be different than its literal meaning.

The problem isn't that there's some other content out there that tells us what is said by a sentence. For example, the Table 8 presents some good candidates in the abstract. The literal meaning of a sentence is always a candidate for what is said by a sentence on a particular occasion. The problem is that there appears to be no way to find what is said by a sentence without recourse to information outside of the sentence; namely, the context in which it is asserted.

[13]I don't mean to take a strong position on whether these paraphrases are the most apt for capturing the literal meaning of a sentence. For example, perhaps breakfast should be understood as whatever meal breaks one's fast (and not strictly speaking as a morning meal). Similarly, I'm assuming that "The X" is treated as a Russellian definite description, which is certainly not required.

Table 8: What is typically said by a Sentence

Sentence	**What is *typically* said**
"I've had breakfast"	I've eaten food before noon today
"You are not going to die"	Your death is not imminent
"It's raining"	It is raining here
"The table is covered with books"	The (or 'a') salient table is covered with books
"Everybody went to Paris"	Every person (relevant to our discussion) has gone to Paris at some point before now
"John has three children"	The number of children John has is exactly three

Semantic Minimalists deny the previous thought. They argue that strictly speaking, what is said by a sentence is its literal meaning. And the literal meaning plays a central role in working out what one means by asserting a sentence. Hence when I assert "I've had breakfast", what I've said is that at some point in the past, I've eaten food before noon. But what I mean (given appropriate circumstances) is that *today*, I've eaten food before noon. On the minimalist's telling, the assertion of the sentence always says "at some point in the past, I've eaten food before noon". It is further pragmatic reasoning that gets us to the content meant, namely "I've eaten food *today* before noon", and (given appropriate circumstances) to the implicature "I don't want food right now". The latter two contents, according to the minimalist are part of what is meant with the sentence (and so should be counted as part of that sentence's pragmatic content), but are not *said* by the assertion of the sentence.

Contextualists—I will explore two kinds below: so-called "moderate contextualists" as well as "radical contextualists"—argue that many sentences do not have literal meanings and even where they do, they are often not *what is said* by the assertion of the sentence. For example, suppose I've *not* had breakfast today, but have at some point in the past. Would I be speaking falsely (if in response to a friend asking me at 11am whether I am hungry) if

I asserted that "I've not had breakfast"? The intuitive verdict is that what is said with this assertion is that I've not had it *today*, and the assertion would only be false if I've had it today.[14]

The crucial difference between contextualists concern the role they see sentence meaning playing in determining what is said by the assertion of a sentence in a particular context. Moderate contextualists think that sentence meaning provides a kind of recipe or formula which tells us how to fill in the missing information in the sentence with the relevant information from the context. For example, indexicals like "I", "here", and "now" can be filled in immediately from information about the assertion (who the speaker is, and where and when she spoke).[15] For example, we can say pretty clearly in advance *what* information is needed to get to a correct verdict concerning what is said in the previous sentences (see Table 9). The moderate contextualist claims that this information can be gathered from the context

Table 9: Information needed to determine what is said

Sentence	**Information Needed**
"I've had breakfast"	When? At all?
"You are not going to die"	For how long? At all?
"It's raining"	Where & when
"The table is covered with books"	Which table?
"Everybody went to Paris"	Everyone among which group of people?
"John has three children"	Whether it is "at least" or "exactly"

of utterance in a straightforward manner (we know what information is missing from these sentences, so we know what to look for). The radical contextualist, by contrast, agrees that

[14] We can ask similar questions concerning the assertion "I've had breakfast" when asked the same question in the same conditions (i.e., I've not eaten yet today, but I'm not hungry). Of course, this can become muddled quickly, because I might intend a certain meaning with the sentence, namely the literal meaning, understanding full well that my interlocutor will understand *what is said* as concerning *today*. Is this a case of speaking falsely or merely deceptively? I'll return to this later.

[15] Even minimalists grant that these sorts of expressions are context-sensitive (information that can be immediately and unambiguously ascertained from the act of assertion, and which require no further reasoning to ascertain can have a role to play in determining the literal meaning of a sentence and hence what is said by its assertion.

"what is said" cannot be determined by sentence meaning alone, but that such information cannot be specified in advance. It might not be clear exactly what information we need (in each context). Further, it may be that contextual information not only fills in missing information, but influences what sort of information and the way in which that information determines what is said by a sentence on a particular occasion. Since this phenomenon is ubiquitous, the radical contextualist also holds that there is no explicable relationship between sentence meaning and what is said by a sentence.

To summarize, what is at issue is this: (i) whether, in resolving what one says and means with an assertion, there is such a thing as "what is said" or whether the first notion of content we can safely resolve is closer to "what is meant", strictly speaking; (ii) supposing that there is such a notion, whether "what is said" is the literal meaning of a sentence (and whether this notion is itself intelligible).[16] What is presupposed as at issue in (i) is whether we can specify in advance the significance of a consideration to determining what is said by an assertion. In the following sections I will explore this debate in more detail and then explain my own contribution to it.

Besides inserting my own view into this debate (which is the central contribution of this section), I also try to reconceptualize the debate. Each camp in this debate conceptualizes the landscape differently. Radical contextualists think that moderate contextualists and minimalists are opposed to them insofar as they both think that there is a straightforward way of articulating the connection between sentence meaning and what is said by a sentence on a particular occasion (though they disagree about what counts as a "straightforward way" of doing so). Semantic minimalists understand contextualists (moderate and radical) as denying their position by asserting that we cannot ascertain the semantic content expressed by a sentence (what is said by that sentence) without significant contextual enrichment. Moderate contextualists understand radical contextualists and semantic minimalists as *agreeing* that very few words make contextually predictable contributions to what is said by a sentence on a particular occasion: they have a small disagreement (namely whether there is a small island of semantic content that remains—literal meaning), but otherwise think that

[16]It's perhaps telling how I've divided these questions, i.e. (i) is meant to distinguish radical contextualists from both moderate contextualists and minimalists. Minimalists might rather start with (ii) in order to lump radical and moderate contextualists together.

basically all contextual enrichment of the sort that interests moderate contextualists is to be accounted for pragmatically (if at all).

I try to show that all of these camps actually agree on a central presupposition: *if a sentence has a semantic content on a particular occasion, there must be necessary and sufficient conditions for determining it.* Since all of the views surveyed take a truth conditional view of meaning, according to which the truth conditions of a sentence tell us something about which objects have which properties, it is unsurprising that we get such a view.[17]

4.1.1 Is what is said always said literally?

Semantic minimalism is the view that "what is said" by a sentence is that sentence's "literal meaning": at least according to how I have been employing these terms. Semantic Minimalists do not object to this characterization, though the distinction between "what is said" and a sentence's "literal meaning" is often employed to deny semantic minimalism (as I'll explain below). Cappelen and Lepore (2008, p.143) characterize Semantic Minimalism as follows:

> "The idea motivating Semantic Minimalism is simple and obvious: The semantic content of a sentence S is the content that all utterances of S share. It is the content that all utterances of S express no matter how different their contexts of utterance are. It is also the content that can be grasped and reported by someone who is ignorant about the relevant characteristics of the context in which an utterance of S took place. The minimal proposition cannot be characterized *completely* independently of the context of utterance. Semantic Minimalism recognizes a small subset of expressions that interact with contexts of utterance in privileged ways; we call these the *genuinely context sensitive expressions*. When such an expression occurs in a sentence S, all competent speakers know that they need to know something about the context of utterance in order to grasp the proposition semantically expressed by that utterance of S, and to recognize the truth conditions of its utterance. These context sensitive expressions exhaust the extent of contextual influence on semantic content."

The basic idea is that the semantic content of a sentence—what would be said by that sentence—can be ascertained without recourse to the context of utterance. So that the sentence expresses the same content in all contexts. There is a small selection of words

[17]Semantic minimalists actually think that an advantage of their view is that it lends itself to a tidier metaphysics. After all "properties" aren't the sort of thing that are context sensitive. Though I'm not sure why they think this. At any rate, I've tried to relegate metaphysical remarks to footnotes.

that *do* require contextual enrichment to reach a content, but these expressions can be characterized entirely by their grammatical category; further, the kind of enrichment they involve is trivial to describe. These expressions are:[18]

> "The personal pronouns 'I,' 'you,' 'he,' 'she,' 'it' in their various cases and number (e.g., singular, plural, nominative, accusative, genitive forms), the demonstrative pronouns 'that' and 'this' in their various cases and number, the adverbs 'here,' 'there,' 'now,' 'today,' 'yesterday,' 'tomorrow,' 'ago' (as in 'He left two days ago'), 'hence(forth)' (as in 'There will be no talking henceforth'), and the adjectives 'actual' and 'present.' Words and aspects of words that indicate tense also have their reference so determined."

Contextualists employ a number of arguments against this view. The arguments rely on example sentences which seem to lack a literal meaning or whose literal meaning is different than what is often (or typically) expressed by that sentence. There are roughly four kinds of phenomena that are used to make arguments for contextualism.[19] This division is not necessary disjoint (and probably not exhaustive) of expressions which require contextual enrichment, but each category displays some unity in how extra-sentential information might influence what is said by a sentence:

1. **Context-Shifting Expressions:** These are expressions *outside of the basic set* whose exact meaning seems to shift between contexts. For example, most gradable adjectives such as "tall', "cold", "small", etc. Generally speaking if an adjective can sensibly take a comparative or superlative form, it is gradable. These are said to be context sensitive because a sentence such as "Tim is tall", might be spoken truly when Tim is thought of as a member of the philosophy faculty and we compare him with his colleagues, but the same sentence could be spoken falsely if spoken of Tim in his heyday as a basketball player.[20]
2. **Incompleteness:** There are also sentences which don't seem to express any semantic content (i.e. nothing can be said to be said by them) without some kind of contextual enrichment. Typical of these are adjectives and nouns which take (optional) complements

[18](Cappelen and Lepore, 2008, p.144). Cappelen and Lepore also refer us to Kaplan (1979). I claim that the enrichment is "trivial", but this is actually far from true. The way that demonstratives (or disambiguation) determine *what is said* is actually far from clear.

[19]This division and the examples employed relies heavily on Borg (2007); Braun (2017).

[20]Since gradable adjectives can also often be *vague*, there is a large overlap with the literature on vagueness. See Kennedy (2007, 2019).

such as "ready", "late", "suitable", "enough", "neighbor", etc. The sentence "Tim is ready" cannot be said to say anything until we resolve *what Tim is ready for*. The contextualist takes this information to be supplied contextually: the context makes it clear what activity Tim is ready for and thus resolves the semantic content of the sentence.[21] It might be difficult to see the unity underlying this list (if there is such unity) since a number of relational terms don't seem to generate incomplete contents. For example, the sentence "John is a father" could be completed with "... of Connor", but we can easily imagine many contexts where such completion isn't required for us to evaluate the content of the sentence (i.e. we only need to evaluate whether John is a man who has children or not).[22]

3. **Different Literal Meaning:** There are sentences which have a clear literal meaning, but what is said by those sentences is very rarely that (nor is the literal meaning typical). These sentences often involve some kind of quantification or involve issues with tense (which can also be cashed out quantificationally). Many of the examples I used earlier are of this form. For example "Everyone went to Paris", rarely is used to say that *literally* every person went to Paris.[23] Instead, there is often some contextually salient class of

[21]See Partee (1989); Bach (2005); Gauker (2012) for contextualists of this type.

[22]Worse still, there seems to be a difference between words which take an optional complement and words which are polysemous and have one sense (when without complement) and a different sense (when with complement). For example, consider "sick" vs. "sick of" and "proud" vs. "proud of". Predicating the former has as a consequence that someone is ill (in the first case) or arrogant (in the second case). But these consequences no longer go through when the complement is added. This is genuinely different from a case like "parent" which seems to genuinely admit an optional complement (rather than being polysemous between a lexeme with and one without a complement). See Gillon (2015) for a more in-depth treatment and further references.

A further problem when accounting for natural language data is that sentences which seem equivalent (truth-conditionally) might be judged to be incomplete in one case, but not another. For example, compare "Sam is a neighbor" to "Sam has a neighbor".

Supposing a clean way of distinguishing words with optional complements, we still lack a way of distinguishing among them words which generate incomplete contents. It may be that such facts are at least partially lexical (rather than purely grammatical).

[23]NB: all the contextualist needs is that there are contexts where "Everyone went to Paris" is used *to say* something different than "literally everyone went to Paris". Contextualists will often use this as a test: does adding "literally" to a sentence change what is said by that sentence.

I am avoiding discussion of relevant theorists in this section (reference above, e.g. Sperber and Wilson (1995)), but their favored examples involve sentences that don't just seem to say something different than their literal meaning, but whose literal meaning is close to nonsensical. For example, a waiter, in saying of a customer (who has ordered a ham sandwich) that he is rude, might say: "The ham sandwich at table 5 is a real jerk", without *saying* anything about the sandwich itself (despite what the sentence seems to say *literally*). Or, suppose in response to the question: "Are you the green Honda out front?" (itself a nonsensical question taken literally), one answers: "No, I'm the red Ferrari." A minimalist can stomp his people who went to Paris. The context supplies that class of people.

Likewise if a doctor says "You're not going to die", when you are sick, we often think that what the doctor has said is that you will survive this injury, you will live (for some contextually salient amount of time). The context supplies that the doctor says you aren't going to succumb to your illness.

4. **Indeterminacy:** Finally, some sentences do not have determinate semantic contents (and thus the very idea seems questionable) until they are actually asserted in a context. According to this view, there is no way in principle of accounting for how sentences receive the contents they do on different occasions. The same words expressing the same meanings (i.e. controlling for ambiguity/polysemy/etc.) can be used to say quite diverse things: OR, what amounts to the same, can be true on one occasion and false on another. For example, the sentence "The kettle is black" might on one occasion be spoken truly because the kettle is covered in soot, which makes it appear black. But on another occasion, the same sentence might be false if what is at issue is the color of the kettle itself, then the fact that it is a kettle covered in red enamel makes a difference. One might think that this can be resolved by distinguishing its "actual color" (so that the first sentence is actually false, but pragmatically felicitous). But even this can run into difficulty. If the "actual color" is something akin to the unobstructed color of the object's surface, then we can imagine cases where "The kettle is red" would be false since *actually*, the kettle is silver, but covered with a red enamel. Even if we somehow suppose we have a way of determining what the *actual* object is in every case, this can sometimes fail to render the right verdicts. Sometimes color predicates are used to refer to the inside of objects rather than their surface colors, or to the objects as they appear in normal conditions.[24]

foot here and just say such sentences aren't *semantically* absurd, just false, and their falsity is important in understanding the pragmatic reasoning that follows (e.g. to "The customer who ordered the ham sandwich is a real jerk" and "I own the red Ferrari"). But if it turns out I don't own the red Ferrari (or there isn't one parked nearby), I don't think it would count as defense of my lie to observe that only a fool would think that *I* am a car (and if I were a car, would I really be an Italian sports car?).

[24]There are a number of other areas where context-sensitive information seems to arise (or in which the literal meaning differs from what is said). For example, epistemic modals like "know" (DeRose, 1995; Lewis, 1996; Rysiew, 2021), modal and deontic language generally (expressions like "could", "might", and also "ought", "should", "must", etc.) (Lewis, 1973; Kratzer, 2012), to name two.

Minimalists offer a number of arguments against the above. One strategy is to treat the kind of context-sensitivity appealed to as actually a case of ambiguity or polysemy. So, for example, we can deal with 1. and 2. by appealing to the fact that there are a number of different properties expressed by the word "tall' (e.g. tall for a philosophy professor, tall for a basketball player, etc.) or "ready" (ready for school, ready to go to lunch, etc.), and that we need to resolve this ambiguity before we can even ask after the semantic content of a sentence. Against this, it is reasonable to observe that there is a clear difference between actual cases of ambiguity such as "Tim has a bat", which involves two sentences that are homophonous (the animal vs. the piece of equipment), and words like "tall", which though they may contribute truly in one context, and falsely, in another, seem to not turn on a lexical difference.[25]

A more plausible minimalist strategy has the effect of reducing 1. and 2. to 3. For example, when we say that "Tim is tall", we are just saying he is above a certain height (perhaps which objects can truly be said to have this predicate varies according to the object, but there is no contextual variance). Supposing "Tim is tall" (but perhaps not taller than most basketball players), we can explain that the assertion that "Tim is tall" (when considering his heyday) as strictly speaking *true*, but infelicitous. It is infelicitous because it is also true that "basketball players are tall" and since it is a pragmatic maxim to be expedient, we think that redundant information is typically used to *mean* something else or something further.[26] This seems less likely to work with sentences like "Tim is ready"[27] and

[25]A related strategy treats such locutions as elliptical. This has some plausibility for locutions that take complements (such as "ready", "late", etc.), but seems rather implausible for gradable adjectives generally. In general these two strategies offload some of the work to "speaker intention", i.e. resolving what the content of a particular expression is pre-semantically (before semantic contents can be assigned).

[26]Another strategy, perhaps close to the one in the previous paragraph, though which allows a single literal meaning to be reached is to postulate a hidden existential quantifier. "Tim is tall" means "there is some height, above which Tim is"; "Tim is ready" means "there is some activity, for which Tim is ready"; "the kettle is black" means "there is some way, in which the kettle is black". This approach runs into the problem that these sentences don't seem to be saying anything about the existence of such standards. "Tim is tall" seems to *presuppose* some standard, but presuppositions aren't part of the semantic content of a sentence.

Even if this strategy worked, it would have the same effect as that which this paragraph is elaborating: that 1. and 2. (and 4.) reduce to 3.

[27]Then again, perhaps there's just some psychological state of "readiness". My intuitions seem to diverge when considering sentences such as "Tim is happy/sad/angry" even though Tim might be happy/sad/angry about different sorts of things. Suppose Tim's father died on the same day that Cubs won the World Series and that Tim is a Cubs fan. Tim's friend are discussing how Tim is doing in light of his father's death. It would seem false to report: "Tim is happy" (even though he is quite happy that the Cubs won).

at least slightly problematic in the cases involving 4.[28]

Regardless, what the above strategies get us is a way of assigning a literal meaning to such sentences without recourse to the context of utterance. That is to say, I believe the Semantic Minimalist has all the resources she needs to explain away 1. and 2. The above doesn't work, however, for 4. The point of the sort of phenomena explained in 4. is that what we take to be the *literal meaning* of the expression is under-determined by the sentence, and indeed, how to fill in the details is likewise under-determined. If the question is of the appearance of the kettle, then we can quite literally say truly "The kettle is black". If the question concerns a more abstract conception of the object (supposing we cleaned it up), then we can quite literally say its negation truly "The kettle is not black" (that is, "the kettle is black" would be said falsely).

Let's pause to take stock. I've elaborated four sorts of examples of expressions that are typically given to challenge Semantic Minimalism (the view that what is said by a sentence is its literal meaning). I've also explained that the minimalist has the resources to give a literal meaning for 1. and 2., so that what remains to challenge minimalism is 3. and 4.

One strategy which is available to both the minimalist and the moderate contextualist is to stamp one's feet on 4. There just is a literal meaning of a sentence like "the kettle is black", or at least a limited range of available options. If they stamp their feet, then the dispute between the minimalist and moderate contextualist concerns the validity of 3. Of course they just have to stamp different feet. The minimalist must either insist that sentences like "the kettle is black" are ambiguous, or that there's just one literal meaning.[29] The moderate contextualist, on the other hand, can say that it varies across context. Exactly how is perhaps slightly more complicated to describe (when compared with some of the other phenomena surveyed: it certainly isn't predictable from grammar how an expression like "... is black" will be contextually enriched). What both moderate contextualists and minimalists agree on is that sentences have semantic contents because there is always a way of determining

[28] In 4. it isn't even clear what the candidate would be in a large number of cases.

[29] I suspect the minimalist will plop for the "actual color" in color cases, but basically any predicate can be subject to this kind of treatment. For example, another Travis-case involves the sentence "Zoe is at home". If the question is whether Zoe would answer the phone should we call, then it makes a difference whether she is inside her house or in the garden. If it concerns whether we could find her by going to her address, then she can be at home (but quite literally outside the physical structure). If the question concerns whether Zoe is traveling or not, she could be two towns over at her friend's house and still truly be said to be "at home".

that content. The disagreement concerns how much contextual information is allowed to influence that determination (and what sort of contextual information). Both think that the grammar of a sentence tells us how to find a truth value of a sentence. The difference is what sort of information the grammar tells us to consult. For the mininmalist—with a few minor exceptions: e.g. speaker, time of utterance, etc.—the only additional information needed to be able to determine a truth value is found in the lexicon. For the contextualist, a lot of contextual information may be allowed in, though what this extra information is, is similarly predictable from the grammar of the sentence (plus some lexical facts).[30] As a result of this, we can say that both minimalists and moderate contextualist agree that:

If a sentence has a semantic content, then the recipe for ascertaining that content must be determinable from the grammar of the sentence (plus relevant information from the lexicon).

One strategy which is available to both the minimalist and radical contextualist is to deny the intelligibility of 3. What is said, strictly speaking, by a sentence is its literal meaning. Sometimes we mean something else with that sentence's assertion and sometimes what we mean, *strictly speaking*, can be quite close to the literal meaning. But *what we mean* is not part of the semantic content of the sentence. If they both insist that "what is said" (as it is distinguished from "literal meaning") is a pragmatic notion, then the dispute between the minimalist and the radical contextualist concerns whether there are sentences whose literal meaning is not determinate (i.e. whether there are genuine cases of 4.). Minimalists take such contents to be determinable, and hence for such sentences to have content. Radical contextualists take them to be indeterminate and hence to lack semantic content. As a result of this, we can say that both minimalists and radical contextualists agree that:[31]

If a sentence has a semantic content, then the recipe for ascertaining that content must be determinable from the grammar of the sentence (plus relevant information from the lexicon).

[30]For example, it might be a lexical rather than grammatical fact that tells us which predicates are context-sensitive or generate incomplete contents without pragmatic enrichment.

[31]Compare Travis (2008, p. 152): "But of course I do not think that, in that sense [viz. predictable as a function of some set of parameters], any sentence does have a truth condition."

4.1.2 The Common Presupposition

In the previous section I argued that the primary participants in this debate seem to agree on a central claim: if a sentence is to have a semantic content, then we need a way of determining what that content is from something like the grammar of the sentence. The grammar of the sentence should give us either the content immediately (after consulting the lexicon), or tell us which contextual information we would need to supply a content. Radical contextualists deny the consequent and thus the antecedent. Both moderate contextualists and minimalists affirm the antecedent (and thus the consequent) but disagree about what sorts of contextual enrichment is licit.

In this section I make two moves from the ground currently covered. I don't think the moves are controversial, but I also can't offer the soundest of arguments for them, so I won't really try. What I'm trying to convey is a way of understanding this debate. And once you've understood it this way, I hope, it will easier to see how the account of substructural content I've been advancing can be used to fill a lacuna.

The first move is to fill in the details of the agreed upon presupposition. Since by semantic content is meant "truth conditions", we should understand a "recipe for ascertaining that content" as some way of moving from a sentence to truth conditions. While there is much that could be said here, the main thing we need is the existence of function for each sentence which maps contexts (indices, to be precise) to truth values.[32] Since indices are of finite size,[33] this means there must be a list of considerations (who the speaker is, when and where she spoke, object of deixis (if applicable), relevant class of tall people, activity under discussion (for which one could be ready), etc.). To sum this up, there is a finite number of considerations needed to evaluate the truth of any particular sentence. These considerations should be *entirely* predictable from the grammar of the sentence itself (perhaps with some consultation of the lexicon). This doesn't mean that a competent speaker could list them, just that what sort of information could be relevant to the evaluation of a sentence is predictable entirely from the grammar of the sentence itself.

[32] I am being rather imprecise here. We should say that each part of a sentence (potentially) receives an index. So that an entire sentence could be said to receive a tuple of indices or something like this.

[33] Actually what seems important is that the values be recursively enumerable. What what we want is that the index attaching to each term be finite.

The second move is to say that because we can evaluate the truth of a sentence given a finite list of considerations that we can of course generate a list of necessary and sufficient conditions for the truth of that sentence. In a certain sense this move may seem a bit controversial. After all, there is a potentially infinite list of speakers. I do not wish to claim that (holding everything else fixed) we have a way of listing all speakers and assigning truth values (the way we map speaker to truth value in such a case may depend upon some further function).[34]

Instead what I want to claim is this. Take a sentence A. There is list of considerations $\gamma_1, \gamma_2, \ldots, \gamma_n$ (some of which could be disjunctive) which are individually necessary and jointly sufficient. In symbols:

$$\begin{aligned} \gamma_1, \ldots, \gamma_n &\vdash A \\ A &\vdash \gamma_1, \\ &\vdots \\ A &\vdash \gamma_n. \end{aligned}$$

For simplicity, I may write $\Gamma \vdash A$ and $A \vdash \gamma_i$ for $1 \leq i \leq n$. These exhaust the truth conditions of the sentence. These conditions are clearly structural.

Monotonicity: Since these are the only considerations relevant to the truth of the sentence A, no further consideration can infirm the implication. Thus $\Gamma, \Delta \vdash A$. Likewise, given that A holds the presence of further considerations may have some effect on the context, but not on the necessary conditions of A, thus $\Delta, A \vdash \gamma_i$.

Transitivity: Suppose $A \vdash B$ and $B \vdash C$. Let Γ_A be shorthand for the considerations which are jointly sufficient and individually necessary for the truth of A. Clearly this means $\Gamma_C \subseteq \Gamma_B \subseteq \Gamma_A$. Since '$\subseteq$' is transitive, so is $\vdash$.[35]

[34] To be precise, I am ignoring the effect that elements of the "basic set" have on content. Everyone in the debate seems to agree on how these work. The question is whether there is a notion of content once those elements of a sentence have been properly resolved.

[35] To avoid a complicated metaphysical debate, I avoid talking about contraction and reflexivity in this chapter. The same is true of particularism below: to avoid the same sort of metaphysical debate (though there concerning action rather than states of affairs).

Of course this is exactly what the radical contextualist denies is possible. The radical contextualist's claim isn't that a sentence can't be true on particular occasions (i.e. given certain considerations). Rather, the claim is that further considerations can affect not just the truth value, but what considerations would make the assertion of the same sentence true. To put the thought simply, it may be that: $\Gamma \vdash A$ but $\Gamma, B \not\vdash A$ since the presence of B makes it such that A is no longer true given Γ. This means that the considerations which make A true do not do so monotonically.

For a similar reason, we do not get transitivity of implication for such sentences. Supposing $A \vdash B$ and $B \vdash C$, just because $\Gamma_A \vdash A$ does not mean that $\Gamma_A \vdash B$ even if $\Gamma_A \supseteq \Gamma_B$. Since, it may be that Γ_A contains within it considerations which change which considerations would make the assertion of B true. More importantly, if $\Gamma \vdash B$—i.e. the collection of considerations, Γ, would make the assertion of B true—it does not follow that $B \vdash \gamma_i$ (for $1 \leq i \leq n$). This is because even though a set of considerations might be sufficient for B, none of them may be necessary (or even any combination of them).

It seems clear then that all participants in the debate are committed to the idea that if a sentence is to have a semantic content then there must be necessary and sufficient conditions for its truth (i.e. some accounting of considerations for its truth). This idea in turn is wrapped up in the thought that semantic content must be structural (i.e. monotonic, transitive, etc.). As I've been arguing, however, we can reject this shared assumption. This allows us to maintain with the radical contextualist that the considerations which would make the assertion of a sentence true on a particular occasion may vary radically from occasion to occasion: that is, that there may be no account of how considerations may combine to yield truth evaluations of sentences. Nevertheless, we may maintain with the minimalist that sentences have stable semantic contents. We arrive at such a view by understanding the meaning of a sentence in terms of its contribution to good implication. This view allows us to have exactly the sort of patterns of meaning that seemed problematic above.

4.1.3 Answering a Panoply of Objections and Other Concerns

In this section I see whether my view stands up to the standard sorts of tests and objections present in the literature. What I mean is, can it account for the sorts of items in 1.–4., and am I able to respond to the sorts of more fundamental objections given by minimalists (which I shall detail below). Since I am attempting to carve out new logical space in the debate, it might seem unclear why I am answering some of the questions as I am (sometimes I seem to want to back the minimalist; sometimes the contextualist). To be clear, what I want to do is maintain the following sorts of intuitions:

- One and the same sentence (i.e. expressing the same meaning) may be the consequence of (and has as its consequences) quite diverse things on different occassions (in my idiom: the same sentence may imply and be implied by diverse things in diverse contexts).
- With a handful of exceptions (predictable by grammatical category) words do not change their contents across contexts.

In order to maintain these intuitions I aim to:

- Account for the purportedly context-sensitive phenomena in a way that is satisfactory to the contextualist (no foot stomping or postulation of suspect literal meanings).
- Nevertheless, show that most of this purportedly context-sensitive phenomena can be truly understand to fail the tests that minimalists advance for context-sensitivity.[36]
- Finally, I explain how my notion of substructural content doesn't face the sorts of objections the minimalist raises against the contextualist. My responses differ from the minimalist in important respects, so it is worth going through them.

4.1.3.1 The Phenomena

The goal of this section is to show that my view of content can accommodate the sorts of cases typically invoked by contextualists and accommodate them in a way that is at least facially in line with what minimalists say, i.e. that these sentences have the same content across contexts. In my own view, I see nothing wrong with allowing for a wide

[36]NB: I don't endorse these tests as tests of context-sensitivity, but I'm willing to spot the minimalist this claim for the sake of debate.

variety of context sensitive expressions. But the goal of the present work is to investigate a particular notion of content—one which allows for substructural implication to be built in at a fundamental level—and see how far that takes us. That is, to see how much of the work taken to be done by appeals to context-sensitivity, can instead be handled in a straightforward way by appeal to the substructural character of consequence. I think such a notion of content can take us quite far.

1. Context-Shifting Expressions: My claim is that the sentence "Tim is tall" requires no further contextual enrichment to receive a semantic content. Certain sets of considerations (i.e., those that involve Tim as a member of the philosophy department) may imply this sentence, while another set (i.e., those that involve Tim as a basketball player may fail to):

$$\Gamma \vdash t,$$
$$\Gamma' \nvdash t.$$

Unlike the previous accounts considered, we can also acknowledge that further considerations in both cases may flip the polarity (perhaps Tim is a member of a particularly tall department; or the question concerns Tim as a member of the intramural faculty team).

Likewise, what follows from this sentence may vary. So when we consider a sentence such as "Tim is over 6'5" tall" it would follow if the considerations are of Tim as a basketball player (and fail to follow as a member of the department):[37]

$$\Gamma, t \nvdash h,$$
$$\Gamma', t \vdash h.$$

2. Incompleteness: The claim here is nearly identical as with the previous one. The main difference between 1. and 2. concerns how they are typically accounted for, but I believe they can be accounted for identically. Perhaps considerations Γ imply that A (meaning "Tim

[37] I am being somewhat inexact here. Γ is being used to mean both considerations that imply t and those that merely make a certain precisification of t more salient.

is ready" (in the sense of, "... for the exam") and Δ in the sense of "... for his trip", then we have:

$$\Gamma \vdash A,$$

$$\Delta \vdash A.$$

But if being ready for the exam means "Tim will pass the exam" (p) and ready for her trip means "She has packed weather appropriate clothing" (c), then we shouldn't want either implication to go through:

$$A \nvdash p,$$

$$A \nvdash c.$$

Since the fact that "Tim is ready" is insufficient for either. Though of course, there may be some further contextual information (together with A that gives us p (or c). Most easy would be Γ and Δ, respectively:[38]

$$\Gamma, A \vdash p,$$

$$\Delta, A \vdash c.$$

3. Different Literal Meanings: Since *what is said* by a sentence doesn't vary, the idea that a sentence's "literal meaning" could vary from *what is said* by the assertion of that sentence, is not intelligible. Nevertheless, while I agree with the minimalists that a sentence should receive the same content in all contexts (for me: contribution to good implication), I also think the contextualist is right that the literal meaning of a sentence—where this is understood as a theoretical construct—is an artificial restriction on the meaning of a sentence, and thus might (and often does) miss many of the uses to which a sentence is put. Let me explain further.

In many cases, I do think that we can isolate what is said, "strictly speaking", by a sentence. When I use the term, I am referring to isolating the *fully structural* contributions

[38]Neither of these further implications give us $\Gamma \vdash p$ nor $\Delta \vdash c$ absent transitivity, which I am assuming we need not impose.

In line with my critique of contraction above, it should also be straightforward to see how examples of this kind *could* be used to generate violations of contraction (but I don't pursue that here).

that the sentence makes to good implication. Thus those contributions which are monotonic and transitive.[39] In chapter one, in explaining the well-behavedness of my notion of content, I explained that we could isolate such a content and even introduce a sentential operator for marking such sentences. As I remarked then, such an operator might be interpreted as "literally" or "strictly speaking" (if it is meant to mark *all structural features*).[40]

Sometimes we can do this in total abstraction (i.e. absent any further considerations). Sometimes a sentence can only be assigned a literal meaning once certain further considerations are accounted for; likewise the same sentence might receive divergent literal meanings given divergent considerations. Let me quickly provide a few examples of each sort of thing.

Consider a sentence such as "Smith is not the nicest guy", which is often taken to *say* that Smith is unpleasant, but according to standard Gricean analysis (and to the minimalist) all the sentence says *literally* is that there is someone nicer than Smith.[41] Given that "there is someone nicer than Smith" is a distinct sentence (it has distinct consequences and is implied by distinct premises)[42,43] it will form a part of the sentence's literal meaning. The example is meant to show that (apart from those guaranteed by e.g. classical logic) a sentence may have as part of its contribution to good implication, a part which is fully structural.

---

[39] As well as contractive and reflexive.

[40] As I remarked earlier. The contextualist must interpret such phrases as serving a purely pragmatic role. They don't change what we say, just cancel most implicatures.

[41] This is, I hope, the least controversial cashing out. If "the nicest guy" requires quantification (i.e. there exists someone who is nicer than everyone else... or (if we don't require uniqueness) anyone who is at least as nice as anyone else...). We can already see how multiple literal meanings might be possible from this sentence.

[42] Even truth-conditionally, we should see a difference. It seems possible to assert truly "S is not the p-est n" even though "There exists a p-er n than S" cannot be said truly. Suppose I say "Barry is not the humblest Bee Gee" (the Bee Gees were a disco trio of brothers; Barry is the sole surviving member, he was also the front-man. I don't know if he is humble, so I mean no offense, just to invoke the usual tropes concerning front-men of famous bands). Unbeknownst to me, however, the other two Bee Gees have passed. So the sentence "There exists a Bee Gee humbler than Barry" is false. One might think this counts against it counting as a part of the literal meaning (since it could be false while the sentence is true). But on the notion of content I am advancing, implication is not understood in terms of truth preservation.

At any rate, it should be emphasized that the literal meaning isn't to be identified with a particular sentence (though perhaps sometimes there is such a sentence), but rather with a contribution to good implication—i.e. a set of inferential roles—one which is a (perhaps proper) part of the original's.

[43] Alternatively, we might maintain that sentences do have literal meanings, but the contribution that such a sentence makes to good implication is fully-classical (i.e. the only sentence it implies is itself). This sort of view is most closely aligned with radical contextualists like Travis. Since, of course, we can always say what is said by "p" on every occasion, namely p. If we suppose that every sentence behaves classically, then we could simply identify its literal meaning that way—but I'd like to at least countenance the idea that sentences may have informative literal meanings.

4. Indeterminacy: Finally, there are sentences which seem to have no content (let alone a literal meaning) without contextual enrichment and whose truth conditional content could be further altered with additional contextual information.[44] Given the notion of content laid in this dissertation, it's entirely possible for a sentence to have no fully structural implications whatsoever.[45] Nevertheless, we can say that it is the same sentence (and thus has the same content) that is asserted across these contexts, since it is precisely the contribution that that sentence makes to good implication which is its content.

4.1.3.2 Tests

Earlier I described how minimalists interpret purportedly context sensitive language (my aim was just to get the contours of the debate on the table). Part of their argument for this view consists in several tests which context sensitive language fail (typically). Since my claim is that I can accomodate the minimalist (i.e. that understanding meaning in terms of contribution to good implication allows for a fixed content accross contexts), on my view most sentences should pass these tests. I aim to explain how.

Test 1: Context Sensitive Expressions Block Inter-Contextual Disquotation:[46] If a sentence S contains a context sensitive expression, then the inference from "Sean said 'S"' to "Sean said that S" need not hold. For example, the inference from "Sean said, "My name is Sean" to "Sean said that my name is Sean" clearly doesn't hold since if I made this inference it would clearly have a false conclusion (he never said it and is speaking falsely).

By contrast the inference from "Sean said "Grass is green" to "Sean said that the grass is green" seems to go through without difficulty.

It's worth pausing briefly to examine what's happening in this test when a sentence is context-sensitive (and what fails to happen for context-insensitive language—at least according to the minimalist). Sentences of the form

[44] Though note the possibility explained in footnote 43.

[45] We can even allow that a sentence be non-classical if, for example, it violates reflexivity. In this case, it would truly be indeterminate (and not just indeterminate outside the classical context).

[46] Cappelen and Lepore (2008, p. 88).

“s said ‘ϕ”’

are true just in case s uttered the sentence ‘ϕ’. By contrast,

“s said that ϕ,

is true just in case s expressed the proposition expressed by ϕ. If the proposition expressed by ϕ is different when uttered by me than when uttered by s, then the inference won't hold. That is, if there is potentially context-sensitive expressions in ϕ, the inference doesn't go through.

Of course, the inference does go through sometimes, even in the presence of context-sensitive language; for example, if it is clear that the relevant contextual information would be shared. Suppose that Sean said (of me and him), “We went to the movies,” then, given this contextual information, the inference from: “Sean said, ‘we went to the movies”’ to “Sean said that we went to the movies” holds. This is because the same contextual information is present.

But what about in cases involving language that the contextualist claims is contextual and the minimalist claims is not? Let's take the sentence “Tim is tall” (“Tim is ready” is an easier case I think). I agree with the minimalist that this sentence fails this test. That is, the inference from “Sean said ‘Tim is tall”’ to “Sean said that Tim is tall” goes through (all else being equal).[47]

A contextualist will deny this because supposing that Tim is tall for a member of the department, but not as a basketball player, and supposing that the current considerations make Tim's heyday salient, the latter will be false and the former true. That is, we can say truly “Sean said ‘Tim is tall”’ (since Sean uttered that sentence), but not that “Sean said

[47] Of course there are plenty of times where such an inference fails. Perhaps there is further considerations to the effect that Sean often exaggerates, or often lies, or is just really bad at estimating height. In this case the inference would be blocked, but this shows a problem with the test rather than the sentence.

If someone is a known exaggerator, it seems plain that we often don't take them at their word. Because that Craig uttered “p” doesn't mean we should take him to have said p (as those words are said by me).

Again (nothing hinges on the argument on this footnote): suppose that Sam records his conversation with Chris. Sam asks Chris about an event she recently attended. Chris reports that “There must have been a million people there”. Later Sam confronts Chris for lying. He reads a transcript of the conversation back saying “Did you not say that there must have been a million people there?” It seems reasonable for Chris to reply that “I did utter the sentence, “There must have been a million people there,” but I clearly didn't say that there were a million people there, if the latter is meant to attribute to me a precise numerical estimate of the number of people there”.

that Tim is tall"' (since, given current considerations, the latter sentence would be false). Part of what my account is liberating us from is a truth-conditional account not only of semantic content, but of implication as well. Thus, the relevant question isn't whether it's possible for the antecedent to be false and the conclusion to be true, but what it means (what follows) from the assertion that "Sean said 'Tim is tall"', and it follows that "Sean said that Tim is tall". Since—to put it into more familiar terms—given present considerations, the former sentence has as a consequence (say) that Sean implied Tim is over 6'5"—though he in fact didn't.

One might object to the way I've cashed this out. It might be claimed that I'm trying to read the current contextual information back into Sean's assertion. As I've been urging, this way of cashing out the case doesn't make sense. Absent further considerations what one does in saying that Sean said, "Tim is tall" is to ascribe that claim to Sean (and all the implicit commitments it carries with it—given current considerations—to Sean). One might say, for example, suppose that the speaker and listener were both present when Sean said "Tim is tall" and so both know it is true. In such a case, the salient considerations which attach to Sean's assertion attach also in the appropriate way to "Sean said 'Tim is tall"', that is, we could (given these shared considerations) say (truly if one likes) that Sean implied Tim is over 6'5". Similar attempts to get the case to work—for example supposing that we preface our assertion that Sean said 'Tim is tall' with the contextual information present for Sean's claim—face the exact same problem (namely, the presence of any consideration sufficient for the former should carry over appropriately to the latter).[48]

Test 2: Context Sensitive Expressions Block Collection Descriptions:[49] This is a variant on a few well known tests for ambiguity.[50] The idea is that if an item is ambiguous then it should resist being collected or used in elliptical constructions.[51] Classical example

[48]Incidentally, it seems that felicitous speech reports always involve reporting the appropriate considerations (so that implicit claims are not misattributed). I also think it's possible that "S said 'p"' (where p occurs in quotations) can be said falsely even though S literally uttered the sentence p if, for example, given considerations, it would be false to attribute p to S.

–S never said *that.* –No, she literally did. –Yes, she uttered that sentence, but we were talking about...

[49]Cappelen and Lepore (2008, p.90).

[50]Zwicky, Arnold M. and Jerrold M. Sadock, 1975. "Ambiguity Tests and How to Fail Them."

[51]Cappelen & Lepore are only interested in the lexical versions of these tests, which are also intended to capture e.g. structural (= syntactic here) ambiguity. The sociological explanation is that most moderate

is "The colors and the feathers are light" (which is supposed to be judged zeugmatic, i.e. competent speakers think there's a problem with an assertion of this). The same is claimed to hold of context-sensitive expressions. So, for example, if it is true (on one occasion) that the "The book is here", and on another that "The TV is here", it surely does not follow (on perhaps some other occasion) that "The book and the TV are here" (since "here" is context sensitive). For comparison (and this is an example of both an unambiguous and context-insensitive expression—or at least intended to be): "The grass is green" and "the leaves are green" seems to allow "The grass and the leaves are green". The minimalists claim that the examples used by contextualists fail this test (i.e. collective descriptions are blocked).

I want to claim with the minimalist that many of the examples invoked by the contextualist fail to block collective descriptions. Nevertheless, I agree with the contextualist that something is inapt if we don't attend to what happens when we try to collect across contexts. I think my own account can help us see past some of the weeds here.

Here's an example: Travis claims that even predicates like "green" are context sensitive owing to the indeterminacy of language. He argues:

> "A story. Pia's Japanese maple is full of russet leaves. Believing that green is the colour of leaves, she paints them. Returning she reports, 'That's better. The leaves are green now.' She speaks truth. A botanist friend then phones, seeking green leaves for a study of green-leaf chemistry. 'The leaves (on my tree) are green,' Pia says. 'You can have those.' But now Pia speaks falsehood." (Travis, 2008, p. 111)

Pia said truly "The leaves are green" and then said falsely "The leaves are green". Can we (or Pia) say truly "The leaves are and are not green"? Let us modify the case slightly. Pia paints her Japanese maple leaves green and reports truly, "The leaves are green". Her friend calls asking for green leaves and Pia remembers that the leaves on her oak tree are green. She says truly, "The leaves (on my oak tree) are green." Clearly she could *not* have said truly, "The leaves (on my Japanese maple tree) are green", so the conditions for "...is green" differ here. Can we collect these two, i.e. to say "The leaves on my Japanese maple and the leaves on my oak tree are green"?

contextualists allow that syntactic ambiguity is settled largely by e.g. speaker intention (since whether a sentence is syntactically ambiguous is a fact that precedes parsing it). Radical contextualists typically don't go deep enough in the details to discuss these kinds of cases.

Clearly if the considerations are those pertaining to her chemist friend, then she cannot, because she cannot say truly "The leaves on my Japanese maple are green". So some considerations allow the conclusion of one sentence, but not the other:

$$\Gamma \vdash t_1,$$

$$\Gamma \nvdash t_2.$$

In this case we fail to get the conjunction:

$$\Gamma \nvdash t_1 \& t_2.$$

We also fail to get the conjunction when mixing considerations, so the considerations relative to the Pia's true assertion regarding her maple and those relative to her true assertion regarding her oak, are different. Hence we don't get their conjunction.[52]

$$\Gamma \vdash t_1.$$

$$\Delta \vdash t_2.$$

We would need monotonicity to ensure that $\Gamma, \Delta \vdash t_1 \& t_2$.

To summarize, the inference from from "The grass is green" and "the leaves are green" to "The grass and the leaves are green" is good, provided the same considerations are present for both of the former. If they are not, then the latter need not follow.

Test 3: Context Sensitive Expressions Pass an Inter-Contextual Disquotational Test and Admit of Real Context Shifting Arguments:[53] Suppose S contains a context sensitive expression. Then it should be possible to assert:

"There are or can be false utterances of 'S' even though S."

[52]If we used a different sort of conjunction we might allow $\Gamma, \Delta \vdash t_1 \& t_2$, but this assumes that the significance of the considerations in Γ and Δ, respectively, can be neatly combined to the considerations in $\Gamma \cup \Delta$, which needn't be the case.

I should also pause to note that if the test is simply whether, given t_1 and t_2, we have the conjunction, then this clearly holds:

$$\Gamma, t_1, t_2 \vdash t_1 \& t_2.$$

Regardless of what Γ is, so long as t_1 and t_2 are classical in the right sort of way.

[53]Cappelen and Lepore (2008, p. 104).

For example, a context sensitive sentence such as "I am in Pittsburgh" could be falsely asserted (by someone not in Pittsburgh), though as it happens, *I* am in Pittsburgh. Thus I can say:

"There are or can be false utterances of "I am in Pittsburgh" even though I am in Pittsburgh.

A context-insensitive sentence like "Grass is green" seems to fail this.

Let us use the tall example again. Are we able to get "There are or can be false utterances of 'Tim is tall' even though Tim is tall"? It seems plausible to deny this outcome (i.e. to agree with the minimalist). Travis, nevertheless, tries to reason that (though cumbersome) we can see how a sentence such as this is context-sensitive (i.e. we don't get the truth of the inter-contextual disquotation):

> "Max could quite correctly say, 'Sid grunts. Of course, one would describe him falsely in those terms (in saying of him "He grunts") if (as might be) what one were thereby saying is that he is in that habit.' Not that I advise so speaking to non-philosophers." (Travis, 2008, p. 157)

Travis' "grunt" example plays on the fact that, among other things, the present tense can be used to report both that something has the capacity to do something and also to report that something does something regularly.[54] But all that Travis has observed here is that it's possible to say (to shift back to my example):

Tim is tall, but it does not follow that he is over 6'5",

which could be truly asserted if the considerations were of Tim as a member of the department. But there's no problem here, because this implication does not follow given current considerations. Notice that Travis does not write that the sentence could be asserted falsely (on another occasion), but that something else *could* follow from the sentence (but here doesn't).

[54]Just about any verb can be made to do this. In walking into a stuffy room, I can ask "Do these windows open" and receive an affirmative answer: "These windows open". I report that my neighbor always opens his windows around 5:00pm and that you can smell freshly cooked pastry. My friend asks, "Do these windows open?" (pointing to a particular set of windows), I answer negatively "These windows don't open" (though they are perfectly capable of being opened. Just suppose the windows are the same in both cases.

What Travis and the minimalist take for granted is this: if something is taken to follow from a sentence and we deny that latter thing, then we must deny the sentence. But we needn't accept that if we don't accept a truth conditional account of content (and thus of implication). That is precisely what I am advocating. Thus, I *deny* the possibility of:

> There are or can be false utterances of 'Tim is tall' even though Tim is tall.

as the minimalist explains we must, while still noting possibility of:

> Tim is tall, but it does not follow that he is over 6'5".

which is precisely the sort of sentence Travis uses to argue against a denial of the former. My point is that these sentences are not equivalent unless one takes a particular view of content. One on which the latter sentence is equivalent to (or at least implies) the former sentence.

4.1.3.3 Further Objection(s)

There are a number of further objections to radical contextualism present in the literature. Following Cappelen and Lepore (2008), I address one.[55]

Radical Contextualism Makes Communication Impossible: It's worth noting that the objection applies to the *radical* contextualist and not the moderate one.[56] It seems that the objection to the moderate contextualist would simply be that it makes communication *difficult*.

The objection goes as follows. According to the contextualist there are a large number of considerations which can change the content (truth conditions) of a sentence. Since (for

[55] Another major objection, which I think straightforwardly does not apply to my account is that:

Radical Contextualism is Internally Inconsistent: The second objections concerns what happens when the tenets of radical contextualism are turned against the view itself. That is, can a sentence like "radical contextualism is true" itself be radically contextually dependent? If so, does that mean that the sentence could be false on particular occasions? In general, how can one espouse a view which claims that the content of claims made may vary wildly from occasion to occasion.

The view that content is substructural doesn't mean that content *always* violates structural features of implication. It only means that it may. So I have two options. Either, (i) this particular claim is a claim that is in fact fully structural. This is the easy way out. Or, (ii) claim about content might vary with context. In a mathematical or logical context, claims about content are much more rigid (since structural features are assumed to hold), but otherwise, claims about content may involve their substructural character.

[56] Cappelen and Lepore (2008) argue that moderate contextualism collapses into radical contextualism. I think this argument doesn't work and the opposite is true (moderate contextualism and minimalism are closer together).

the radical contextualist) there is no way of saying what the limit on these considerations is, it is (absent a miracle) impossible to guarantee that speaker and listener converge on all of these considerations. Since these considerations alter the truth conditions and thus *what is said* by an utterance, it is impossible (absent a miracle) for *what is said* by a speaker to be understood as such by a listener. Since the listener has no access to exactly which considerations are influencing the speaker, she may understand something else to be said by the speaker. But if the speaker cannot be assured *what is said* by his utterance is understood as such by the listener, then the idea that genuine communication happens is impossible.

There are two important points to note surrounding this objection. The first is that it is founded upon a Gricean idealization concerning how communication works: namely that assertion requires that the speaker and listener converge *exactly* on the content of the assertion for it be such an act.[57] The second—in connection with the first—is an agreement (among both minimalists and radical contextualists) that most of the information transmitted during communication is imprecise, messy, or open to revision. What I mean with my assertion can be misinterpreted by my interlocutor. What the minimalist wants (but the radical contextualist claims he cannot have) is a stable foundation on which all this messiness is built. We might not be able to say exactly what a person means with a particular assertion on a particular occasion (since the significance she attaches might involve a host of factors that aren't fully transparent to us), nevertheless we can at least say *what she said.* And it is this stable foundation that lets us coordinate and reason (with at least some reliability) upon what someone means or means to do with their assertion. According to the radical contextualist, we cannot converge upon this, so when we question what someone might mean or presuppose with their assertion, we do so with no guarantee that we even understand the same thing *to be said* by such an assertion.

I do not in general agree with the Gricean idealization and the picture of language it presupposes. Nevertheless, the picture of content I am advancing is capable of meeting this sort of objection. Since the content of a sentence on my view is the contribution it makes to good implication, what a speaker says by a sentence is understood as such by her interlocutor.

[57] Compare Grice's formulation concerning meaning and Lewis on convention.

Nevertheless, it may be possible for the interlocutor and speaker to not understand exactly which consequences are intended by a speaker (if their understanding of what considerations are active differs). But they are perfectly capable of negotiating this, in the moment.

4.1.4 Taking Stock

The debate over radical contextualism and semantic minimalism turns on whether sentences have stable semantic contents or whether their contents may vary (potentially radically so) from occasion to occasion. In this section, I argued that there was a common presupposition underlying this debate, which is that if a sentence is to have a semantic content, then that content requires the specification of necessary and sufficient conditions for the truth of the sentence. Such conditions might make reference to nothing extrinsic to the sentence; or, as the moderate contextualist holds, such conditions might need to appeal to considerations outside of the sentence (but otherwise predictable from the grammar of the sentence and perhaps some lexical facts concerning terms that appear in the sentence). Since all participants granted that the content of assertion is truth conditional, they seem to grant that a requirement for having a semantic content is a specification of necessary and sufficient conditions.

I argued that such a requirement on semantic content forces a structural understanding of content (i.e. one in which the contents of sentences stand in fully structural relations of implication). If we abandon such an understanding, and instead opt for a substructural understanding of content it is possible to both (i) maintain that sentences have stable semantic contents (namely their contribution to good implication) and (ii) that the patterns of substructural implication involved in this content undergird the linguistic intuitions that motivate the contextualist.

4.2 Moral Particularism and Moral Generalism

Now, I'd like to do for the debate concerning moral particularism, much of what I did for the radical contextualism and semantic minimalism debate. That is, I'd like to show how a common presupposition underlying the debate presupposes a notion of content that is fully structural. And that just as in the previous debate, abandoning such a presupposition in favor of a substructural notion of content opens up new logical space in the debate. I'll start by motivating the debate and introducing some key distinctions.

When asked 'why' following a choice to pursue one course of action over another, or preceding such a choice in deliberating about what to do, we often look for reasons to help motivate us or answer this 'why'-question. One-and-the-same thing can often seem to take a number of different forms. For example:

- Lying is wrong.
- Lying is bad.
- One ought not lie.

What all of these amount to is some sort of connection between the recognition (on a particular occasion) that an act ϕ would be a case of lying and that act being wrong/bad/prohibited. We can call such a connection (and since the previous three things are meant to invoke such a connection) a principle. If the principle has to do with rightness/wrongness; or good/bad (in a moral as opposed to prudential or instrumental sense); or an ought/must that is thought to overrule other such oughts/musts, then we can call such a principle a *moral principle*. This recognition (or the thing recognized) is often called a (or the) reason for the relevant normative verdict.[58]

Some think that there is a finite set of principles, or at least some formal characterization of any possible principle, that allows us to say with some degree of abstraction whether or

[58] The literature on reason(s) is so vast, I can't pretend to even have a grasp of its geography, let alone to have surveyed this or that particular peak. Nevertheless, I will say that I don't mean to commit myself to any substantive view concerning what reason(s) are: e.g. states of mind, concrete things in the world, properties, etc. I hope my discussion can stay neutral on such matters as much as possible. While they aren't unimportant to this debate, they won't figure centrally in how I frame things. Actually, as I explain below, the assumption that this thing is the "reason" plays the same role as truth-conditionality played in the previous debate (i.e. it's something I'm going to away from by the end of the section).

not an action would be right or wrong. For example, Kant held that an action shouldn't be performed unless its maxim (rule according to which we perform it) can be willed to be a universal law of nature:

> "Act as if the maxims of your action were to become through your will a universal law of nature." (Kant, 1900, 4:421)

So that we can obtain specific principles such as "Do not lie!" through the above (i.e., the recognition that if an action would be the telling of a lie then it could not through my will become a universal law of nature). Scanlon formulates a general principle as follows:

> "an act is wrong if and only if any principle that permitted it would be one that could reasonably be rejected by people with the motivation just described [viz. "people who were moved to find principles for the general regulation of behavior"] (or, equivalently, if and only if it would be disallowed by any principle that such people could not reasonably reject)." (Scanlon, 1998, p. 4)

Here a principle such as "if lying would be more profitable than telling the truth, I should lie" is a principle that could be reasonably rejected by someone who was moved to find principles for the general regulation of behavior (since it would in general be to our disadvantage if we could not in general rely on the words of others as truth, especially in cases where lying would be advantageous to the speaker).

A problem with this sort of view, however, is that the principles which guide us are rarely ever so straightforward. Yes, it is true, that "lying is wrong" is a moral principle we generally acknowledge, but nevertheless, sometimes we think lying is *not* wrong (or would be the less wrong action to take). For example, if the lie is a particularly small one such that the deception would not have a deleterious effect on the addressee's future choices, and if, in addition, telling the truth would have such an effect or would cause immediate harm to the addressee, then we often think that lying is at least permitted (if not the *right* thing to do). Consequentialists, for example, explain this verdict (that, in that case, lying is *not wrong*) by appeal to the good/harm done as a result of the action. Since, in this case, telling the truth is *bad* (causes harm), lying is permitted (or even right).

A distinctive advantage of the consequentialist approach is that the rightness/wrongness of an action swings free of features intrinsic to the action. It is (to a certain extent) irrelevant whether an action would be a case of lying to the rightness/wrongness of an action. What matters are the harms and benefits that result from the telling of that lie. Many have been uncomfortable with this sort of view. They agree with many of the intuitions generated by the consequentialist—for example, that the telling of a minor lie is permitted—but they want to maintain the idea that, for example, it is *in general* wrong to lie (or break a promise/kill/etc.). That is, they want their ethics to be able to make sense of the idea that "to ϕ would be to lie" is a *reason* against ϕing (since lying is, in general wrong).[59] W.D. Ross, for example, held such principles (and so the duties they generate) are only *prima facie* or *pro tanto* principles. While it is a *prima facie* duty not to lie, further considerations (such as those discussed above) allow for the duty to no longer apply. He writes:

> "I suggest 'prima facie' duty or 'conditional duty' as a brief way of referring to the characteristic (quite distinct from that of being a duty proper) which an act has, in virtue of being of a certain kind (e.g. the keeping of a promise), of being an act which would be a duty proper if it were not at the same time of another kind which is morally significant." (Ross, 1930, p. 19–20)

This sort of view is intended to preserve some of the strongest intuitions that motivated each of the views above. That is, we can say truly that "lying is wrong" (and that this is therefore a reason against ϕ, if to ϕ would be to tell a lie). Nevertheless, we can also acknowledge that on particular occasions further considerations, for example that to tell a lie would cause a small harm and that telling the truth would cause great harm, disarms the original reason. To be precise: we have a prima facie principle—lying is wrong—and a further principle—such as: lying is permitted when telling the truth would cause a greater harm (or perhaps: if the difference in harms/benefits of a particular action is significant, do that which causes the least harm/most good)—then the former, applying only *prima facie* is no reason in view of the latter (which is to say that the action in question is "at the same time of another kind which is morally significant").

Finally, I'd like to introduce one further sort of account to lay alongside the three already surveyed (I will return to these accounts throughout this section). This is the account of default reasoning given by John Horty.[60] According to this account certain features of actions

[59] Rule utilitarianism is one such *consequentialist* attempt to do this. Since the point of this chapter is not to defeat nor endorse consequentialism, and since this dissertation could hardly be considered a dissertation in meta-ethics, I am (hopefully permissibly) glossing over quite a bit of nuance where it isn't relevant to the main line of thought.

[60] The most sophisticated and complete version is given in Horty (2012), but see also his 2001; 2007b;

(or further considerations) may have "default rules" associated with them. For example:[61]

$$\text{to } \phi \text{ is to lie} \mapsto \text{you ought not } \phi.$$

We should understand the default rule as a kind of principle; the antecedent of the rule—when present—functions as a reason. The advantage of Horty's account is that we may have in addition principles such as:

$$\text{to } \phi \text{ causes significantly less harm than to not } \phi \mapsto \text{you ought to } \phi.$$

And these two principles come into conflict, since they entail contradictory recommendations. To resolve this, Horty appeals to a ranking of default rules. In this sort of case, the latter rule is deemed more significant than the former, and so the former rule is ignored. In more complicated cases, there may be a number of different default rules, which are triggered (=antecedent holds) and which come into conflict. We resolve the conflict by finding a stable set of rules (=one which doesn't generate any contradictions) and which prefers those rules which are deemed more significant.[62] Like Ross, Horty finds the sorts of generality contained in the Kantian account to be extremely compelling. That we can say, in general—though here the idiom is "*by default*"—that because an action has a particular feature (e.g. that it would be a case of lying) that that counts against doing the action. Officially: that ϕ is a case of lying *is a reason* against ϕ (and this connection is captured by the relevant principle, i.e. default rule). He also finds (again with Ross) it necessary that we be able to acknowledge the way that further considerations may combine in complicated ways to nullify some of the default generalizations.

An important difference between Horty and Ross, however, is the way in which their accounts treat these tentative reasons. Ross' *prima facie* principles are verdictive. That ϕ is a lie means that one ought not ϕ. When two principles conflict, it renders one of the principles inert. Generally, some set of non-verdictive considerations yield an answer. By contrast, Horty's default rules are not verdictive. Rather, what default logic does is tell us which of the reasons we have in a particular context *are verdictive*. I'll have much more to

2007a. Default logic was introduced by Reiter (1980).

[61] I will use '$\mapsto$' throughout for default rules. We should understand statements like the one below to represent the fact that the premise is a (default) reason for the conclusion.

[62] Officially, all the rules are ordered. Stability means only accepting those default rules which are not conflicted by a rule which is lower on the preference ordering (i.e. by a rule which is deemed more significant).

say about this below.

Let's take stock. In accounting for the moral reasons given for action, some accounts locate features within the act itself to determine whether it is right or wrong; some locate features that may be extrinsic to the act, but which nevertheless provide reasons for or against undertaking the action and combine to yield moral verdicts. Finally, some (like Ross and Horty) try to split the difference. That is, to acknowledge that there are intrinsic features of actions (e.g. that an action is case of lying) that provide reason for or against that action, but also acknowledge that further considerations may infirm that original reason

It is against this backdrop that the debate concerning moral particularism and moral generalism is held. Particularists hold it is impossible to account for all of the ways that considerations or reasons may combine to yield moral verdicts about what to do on particular occasions. For example, we seem to acknowledge a principle that "lying is wrong" and a further principle that "if the difference in harms/benefits of a particular action is significant, do that which causes the least harm/most good". So it may be on a particular occasion lying is the thing to do (since to tell the truth would cause a greater harm). But suppose that on a particular occasion, the friend is question recently learned that many people (even his closest friends) routinely lied to him to avoid hurting his feelings. They meant no harm, the slightest insult, even unintended, often led to despondency in him. In light of this discovery, he asked that you promised him to never tell such a lie. Even though he may be more hurt by the truth than the lie, it seems like your promise provides an *even stronger* reason to tell the truth than to lie to him. But suppose he asks you directly whether you've ever told such a small lie to him before? Intuitions can become muddied.

In light of this sort of phenomena, the particularist claims that there are no moral principles; i.e., there are no considerations which apply to actions (or reasons for action) in all cases. Since any feature of an action (or of a reason for action, or of a consideration relevant to action) could be rendered invalid given some further reason or consideration, and since this can happen in highly complex and unpredictable ways, the particularist claims we cannot form any principles for determining the rightness/wrongness of an action in general. This does not make the particularist a nihilist or skeptic about morality. The particularist thinks it perfectly intelligible that on particular occasions we reach moral verdicts. Further, on such occassions we are able to justify those verdicts by appealing to reasons. In short: the particularist believes that the particularities of an occasion are what determine moral verdicts and given their potential diversity, there are no general principles for determining whether an action is right or wrong.

In the sections that follow I explain in more detail what exactly the particularist is committed to and some of the arguments used to support this stance. I likewise try to explain the sorts of reasons given in favor of generalism: (the negation of particularism, namely) the position that there *are* moral principles, i.e. features of action (or of considerations) which provide reason(s) for/against action.

4.2.1 First Distinction: Contributory Reasons and Overall Reason

In order to articulate the particularist's position more precisely, it will be helpful to put two distinctions on the table. The first distinction concerns contributory reasons and the overall reason. Dancy (2004b, p. 15–16) gives a first gloss as follows:

> "A contributory reason for action is a feature whose presence makes something of a case for acting, but in such a way that the overall case for doing that action can be improved or strengthened by the addition of a second feature playing a similar role. Also, a contributory reason on one side is not necessarily destroyed by the presence of a reason on the other side. This does happen sometimes, I agree, but it is far from the standard case. Contributory reasons are officially reasons capable of doing what they do either alone or in combination with others. But they can combine in peculiar and irregular ways, as we will see. There is no guarantee that the case for doing an action, already made to some extent by the presence of one reason, will be improved by adding a second reason to it. Reasons are like rats, at least to the extent that two rats that are supposedly on the same side may in fact turn and fight among themselves; similarly, the addition of the second reason may make things worse rather than better.
> [...]
> But as well as talking about reasons in this way, we also speak of what there is overall reason to do. There is nothing wrong with that, of course, but it should not delude us into thinking that there are such things as overall reasons in addition to the contributory ones. To talk of what there is overall reason to do (and note that 'reason' in this phrase is not a count noun) is to talk about where the contributory reasons come down—on this side or on that."

When we speak of contributory reasons, we mean those considerations, which *contribute* in some way towards our normative verdicts. By contrast when we speak of overall reason, we mean what all of those considerations tell us *all things considered* or *on the whole*. I have ten reasons to do ϕ and ten reasons not to do ϕ, and half a dozen further considerations concerning ϕ. Somehow I end up ϕing or not and hopefully my (in)action is in line with what I had overall reason to do. That is, how those reasons somehow combine to deliver an overall verdict concerning ϕ. To appreciate the difference, consider what happens when I acquire an eleventh reason to ϕ. The eleventh reason combines with all my previous considerations and reasons somehow to deliver a verdict conerning whether or not to ϕ. But the original verdict, the original overall reason I had to ϕ or not, in a certain sense becomes irrelevent. It is replaced with overall reason (to ϕ or not), but in light of a new combination of reasons.

It is important not to confuse several closely related but subtly different distinctions.[63]

1. In jurisprudence, there is a distinction between something which is *probative* and something which is *dispositive*. Often this is applied to evidence (or facts) of a case. Probative evidence is evidence which has some bearing on the outcome of the case, but which needn't settle the issue. If a man is accused of murder, evidence of a motive would be probative, but not dispositive. Dispositive evidence, by contrast, is something that settles the case. Uncontestable video-taped evidence of the accused committing the murder would be dispositive (it settles the issue).[64]
2. Peremptory vs. enticing reasons.[65] This is a distinction between reasons which, on the one hand give us some motivation for doing something, but which need not close off alternatives and which certainly do not get us all the way to a claim about what we ought to do. That something would be fun, is an *enticing* reason for doing it. Peremptory

[63] I say subtly different, but this is a potentially controversial assertion. Some accounts may want to cash out this distinction in terms of one of the distinctions mentioned here.

A further distinction that is sometimes laid alongside these is between moral and non-moral reasons (sometimes they are taken to be the same), since moral reasons are thought to be decisive in a way that non-moral reasons are not. That is, some think that a non-decisive contributory reason must be a non-moral reason, or atleast seem to line up things in this way.

[64] Though note that even video-taped evidence is not equivalent to a guilty verdict. Despite its dispositive nature, there may be sufficient probative facts mounted in defense (or dispositive considerations against—perhaps a video recording of the deceased making death threats).

Alternatively, one might argue that dispositive reasons are a third sort of thing: neither contributory nor overall. Dancy seems to suggest this possibility. Part of the issues hinges on whether it is so much as possible to have two conflicting dispositive reasons (or if this notion is itself a contradiction). The main point here is that the overall reason/contributory reasons distinction is a distinction not *among* reasons, but between levels of reason.

[65] Dancy (2004a,b). This distinciton is owed to Dancy. The original context is meant to show that the moral/non-moral distinction is orthogonal to the contributory/overall.

reasons, by contrast, always issue oughts and demand our attention. This doesn't mean that just because something is a peremptory issue that it decides the issue (after all there may be peremptory issues for and against a certain course of action). Rather, peremptory reasons have a force that enticing ones do not.

3. Insistent vs. Non-insistent reasons.[66] Insistent reasons are reasons which cannot be ignored. By contrast, we may neglect non-insistent reasons and non-insistent reasons may never ground requirements on us (though they may contribute in this or that way).

All of these distinctions are distinctions *among* reasons. That is, they are distinctions among reasons which may combine to yield an overall verdict about what one must do. Some of these reasons take us all (or most) of the way to such a verdict—and Dancy calls such a contributory reason, in general, a *decisive* reason. But these are still reasons that may combine in view of an overall verdict. A man is contemplating murder. That the action he is contemplating is murder, is decisive/peremptory/insistent. That the action would be cruel does not decide the issue, but it does make the decisive reason that much stronger. By contrast, what he has overall reason to do: not murder, is not affected in the least by the further consideration that the murder would be cruel. To further appreciate the difference, consider that there can be no conflicts among overall reason. If there are genuine moral conflicts, then it may be that you have decisive reason (on a particular occasion) to both ϕ and to not ϕ.[67] But these decisive reasons (together with other contributory reasons) combine to give us overall reason (i.e. a normative verdict) about what to do.

Distinguishing overall reason from decisive reason will be important for appreciating the differences between the accounts surveyed earlier. It makes a difference whether we start with the notion of a contributory reason and use further principles to isolate reasons among them as decisive (and then use that decisive reason to reach an overall verdict), or whether we start with the notion of overall reason and use contributory reasons in order to settle any potential conflicts. Thus, a central question concerns how these two notions relate to

[66] Kagan (1989). This one tracks very closely the moral/non-moral distinction.

[67] Consider Sartre's would-be resistance fighter whose mother has fallen ill. He has decisive reason both to join the resistance and to stay.

If a criteria of moral conflictedness is that there is no verdict concerning what to do, then we would have a case where we have a number of contributory reasons (some of them decisive), but no overall reason.

each other. The language employed suggests that we understand overall reason *in terms of* contributory reasons. That is, that wherever we have overall reason to do something, it is precisely because of how various contributory reasons combine. But we need not understand things in this way. We could for example make sense of contributory reasons in terms of overall reason. Or simply abandon contributory reasons altogether.

In the former case, we need some way of accounting for how various considerations contribute towards our normative verdicts. In this domain mathematical language is not uncommon, i.e. we want to know how considerations (chiefly: reasons) combine.[68] Hence theories of rational choice are examples of accounts which place contributory reasons prior to overall reason.

Such accounts may be more-or-less sophisticated or precise, but they give us some way of understanding how, in the presence of such-and-so reasons or relevant considerations such-and-so a verdict is to be reached. What these accounts consist of are principles. Principles which tell us how to get from some set of considerations to a normative verdict (in general or concerning some particular action). Utilitarianism is a very simple version of such an account. It tells us that a consideration counts as a reason (in the contributory sense) for an action iff it counts in favor of the action maximizing happiness. Further, the weight of a particular reason is directly proportional to the degree to which it maximizes happiness. We arrive at overall verdicts (concerning right/wrong) by determining which possible action would produce the most happiness, such that reasons can be compared and combined in relatively straightforward ways. Horty (2012) provides another more sophisticated account of placing contributory reasons prior to overall reason. A "default reason" is a kind of contributory reason. Further, several reasons may combine to defeat one another or otherwise provide a

[68]It won't be of chief importance to me, but a distinctive advantage of treating the contributory as primary is exactly the possibility of agglomeration. Contractualists (like Scanlon), for example, have trouble accounting for the fact that several reasons on one side may be weightier than one similar reason on another. A common objection takes the following form: two groups of swimmers are stranded on rocks at low-tide. One group consists of five swimmers; the other of one swimmer. I only have enough time to rescue one group before the tide comes in, drowning anyone left out there. If I choose rock A, then anyone on rock B could reasonably reject to such a principle; and vice-versa. Hence it seems my hands are tied. But the obvious thing to do is go for the rock with more swimmers.

Likewise, since neither decision fails to universalize, Kantian deontology seems to allow that you have just as much reason to go for the one rock as for the other, which many find to also be the wrong verdict (NB: this is different from the trolley problem where Kantian deontology forbids us from switching tracks). It's not just a matter of charity that you go for the rock with more swimmers, it is required!

stronger reason.[69] One ought to do an action if a "triggered default" issues that judgment and it is not conflicted by a higher ranking default rule (i.e. no default rule that is more preferable says the opposite).

Before returning to the original distinction, it is worth noting what is going on here. The following two things are *not* contributory reasons. They are *not* decisive reasons or dispositive reasons or any such thing.[70]

- To ϕ would produce the most happiness.
- To ϕ is something I have reason to do and such reason is not conflicted by a higher priority default.

These are statements of *overall reason.* They are particular instances of principles which operate at the overall level. They refer to general features of a situation and not particular considerations.

Returning now to the original distinction, some insist that an account of the overall must be given independently of the contributory. An account of this sort could be used to make sense of the contributory; or, the contributory could be abandoned or not given an important place in our ethical theorizing. There are a number of ways this could go, but what this understanding of the overall entails is the desire to understand normative verdicts concerning actions directly, in terms of features intrinsic to those actions (or as closely as possible). Let me run through three examples.

Kantian Deontology: In Kantian deontology we are given a principle that tells decisively whether an action is obligatory/prohibited/permissible from a description of the action *alone.* We need not appeal to any further considerations to determine verdicts considering action. The need to avoid appeal to further verdicts is not just a feature or happy outcome for Kant, it is, for him, constitutive of the ethical. If ethical verdicts were sensitive to further considerations then the verdicts reached would be *conditional* on those considerations, and hence not *categorical.* Since Kant wants to advance an enlightened

[69] Horty is only able to handle agglomeration in a somewhat clumsy way.

[70] Though we might say that whatever considerations combine to make ϕ produce the most happiness are (as said combination) a decisive reason, or that the unconflicted, triggered default is a decisive reason (at least in the particular context).

ethics based on principles of reason, ethics must be categorical. Hence, for Kant, the ethical is given solely in terms of overall reason.

That ϕ would be a gentle murder has no bearing on the overall verdict prohibiting ϕ. It might make something of a case in favor of ϕ against some other act, ϕ', which would be a cruel murder. But both acts are equally prohibited. The cruelty or gentleness of the murder is irrelevant.

Scanlon's Contractualism: Scanlon thinks that normative verdicts (concerning the morality of action) are constituted by our attitudes/motivations/etc. as follows:

> "When I ask myself what reason the fact that an action would be wrong provides me with not to do it, my answer is that such an action would be one that I could not justify to others on grounds I could expect them to accept. This leads me to describe the subject matter of judgments of right and wrong by saying that they are judgments about what would be permitted by principles that could not reasonably be rejected, by people who were moved to find principles for the general regulation of behavior that others, similarly motivated, could not reasonably reject. In particular, an act is wrong if and only if any principle that permitted it would be one that could reasonably be rejected by people with the motivation just described (or, equivalently, if and only if it would be disallowed by any principle that such people could not reasonably reject)." (Scanlon, 1998, p. 4)

Similarly to Kant, Scanlon holds that we can inspect whether an act is right or wrong (so mandatory/prohibited/permissible) by examining the action itself. Unlike Kant, Scanlon is comfortable with some degree of "outside" influence, but such considerations must be formulated in a way to be applicable to everyone (or in fact to already be so applicable). Kant writes that one cannot lie because willing it to be universal would destroy the idea of telling one another anything. Scanlon doesn't go quite so far. Instead, his view allows that we can lie if, for example, the lie is small and the harm it avoids is comparatively large. Everyone seems to agree that it is better to avoid telling the truth when the information which is withheld/distorted would have no significant effect on a person's conduct or life *and* when disclosing that information would have a significant *negative* effect on a person's conduct or life; that is, this is a principle that no one could reasonably reject.

But that any principle that permits ϕ could be reasonably rejected by someone (so described above) is not a contributory reason against ϕ. Nor is it a decisive reason against ϕ. It is a statement of the fact that we have overall reason to not ϕ (the decisive reason, if there is one to be found, resides in the principles that would be reasonably rejected and why they can be reasonably rejected).

In each of the previous two examples, the contributory is abandoned. Instead we are given principles that apply directly to action. In Kant's view these principles are simple; Scanlon, by contrast, allows that such principles may be quite complex.

Finally, consider Ross' account that I detailed briefly above. Like Kant and Scanlon, Ross wants to give an account of overall reason independently of any account of contributory reasons. But unlike Kant and Scanlon, Ross thinks that there is a place for contributory reasons in his account (more on exactly what this means a little later). The important point is that Ross thinks we can make sense of what it means for a consideration to be a contributory reason by appeal to the overall level.[71] That we have *prima facie* reason not to lie means that *absent any further considerations* if ϕ would be a case of lying, then we ought not ϕ. But we can also say that, in the presence of further considerations (and even in the presence of further reasons), that ϕ is a case of lying gives us *a* reason not to ϕ. This reason derives from the *prima facie* reason we have not to lie. Hence, the overall reason we have to not lie, grounds the contributory reason we have to not lie (even where this contributory reason does not settle the issue because of the presence of further considerations).

Let's take stock of the examples I've been using: Kant and Scanlon both provide an account of the overall level independently of contributory reasons; in fact, they think we can simply do without the contributory (at least in our ethical theorizing). Ross also thinks the overall can be accounted for independently of the contributory, but unlike Kant and Scanlon, leaves a place for the contributory and thinks that we can understand what it means for a consideration to be a contributory reason in terms of overall reason. Horty, various species of consequentialists, and decision theorists, by contrast wish to understand the overall *in terms of the contributory*. That is, they understand what we have overall reason to do in terms of how various contributory reasons combine to yield a normative verdict.

[71] I should caution that my version of Ross is probably something of a caricature. First it is informed above all by Dancy's writing on it. Second, Ross never considers anything like "the contributory". He does seem to grant that that which we would call the contributory, is formed via generalizations from the overall level.

4.2.2 Second Distinction: Atomism and Holism

The second distinction is a distinction between *holism* and *atomism* concerning considerations (in particular, reasons). If one is an atomist about considerations then, the significance of a consideration as—for example, a reason to ϕ (or a reason against ϕing)—is completely insulated from other considerations, change of context, etc. We might think the fact that ϕ would be an unjustified killing is a consideration which counts as a reason against ϕing. Further, because of the nature of this consideration (that ϕ is an unjustified killing), we might think that further considerations or particular occasions can't change the significance of this consideration. This doesn't make the reason decisive, however. That ϕ would hurt Paul counts against ϕ, but it does not settle the issue. The question is whether any further considerations can change the relationship between the consideration that ϕ would hurt Paul, and its speaking against ϕ. If one thinks that *all considerations* are like this—that is, they are immune to revision given further considerations, or do not depend upon further considerations for their significance—then one is an *atomist* concerning considerations.

The clearest example of an explicit endorsement of atomism is an endorsement of what is call *separability* or *the sure-thing principle* in decision theory.[72] This principles says that when weighing reasons, different outcomes can be evaluated independently of one another. This means is that if something provides a reason of a certain strength, it should do so independently of other considerations, in particular other outcomes or actions under consideration. This has a number of highly desirable consequences in decision theory. For example, if I prefer (or have more reason) to A than to B; and likewise for B and C, then I ought to have the same relation between A and C (that A is preferred to C). This property (which should be familiar) is called transitivity. Separability guarantees transitivity of preferences (and so seems to have a similar effect on the reason relation).

Sometimes, however, we think that there are reasons against separability (and so against atomism). A famous example is given by Allais (1953) and generally referred to as Allais' paradox. We are to consider four lotteries (L_1–L_4), where depending upon a number randomly generated (1–100) we will receive some amount of money in return, as described in

[72] See e.g. Broome (1991). This is called separability or independence by Jeffrey (1990); it is called "the sure-thing principle" by Savage (1954).

Table 10.[73] If given a choice between L_1 and L_2 (as lotteries to play), many prefer L_2 to L_1.

Table 10: Allais' Paradox

	1	2–11	12–100
L_1	\$0	\$1.1m	\$1m
L_2	\$1m	\$1m	\$1m

	1	2–11	12–100
L_3	\$0	\$1.1m	\$0
L_4	\$1m	\$1m	\$0

The reason is simple: L_2 is a guaranteed \$1m. The 10% chance of an additional \$100,000 that L_1 provides hardly seems worth the 1% risk of nothing.

If given a choice between L_3 and L_4, many prefer L_3 to L_4 (or are at least indifferent on my setup). The difference between a 89% and 90% chance of nothing is not judged as significant. But this is a violation of separability. This is because the weight given to the 1% chance of getting \$0 with a roll of '1' is valued differently when considering L_1 than when considering L_3. To make it plainer: given that we receive the same amount of money when rolling 12–100 in $L_{1/2}$ and likewise with $L_{3/4}$ we should be able to make a decision without considering those columns. That is, consider the choice as redescribed in Table 11. Given that x is identical in each case, discovering that x is 0 (or \$1m) shouldn't change our

Table 11: Variable Presentation of Allais' Paradox

	1	2–11	12–100
L_5	\$0	\$1.1m	x
L_6	\$1m	\$1m	x

preferences, if separability holds. But our intuitions tell us that it does make a difference.

[73] I've seen the lottery described in various ways. Often the cash prizes make L_1 preferable to L_2 and L_3 preferable to L_4 if one *only values money*, *values it linearly*, and uses a decision procedure that weighs the benefit of an outcome proportionally to its probability of occurring. I've chosen numbers that make the choice (given such preferences) equally preferable. One can get this feature back by either making \$1.1m smaller (though larger than \$1m, to avoid the latter lotteries dominating), and/or making a corresponding change to the number ranges.

That an additional consideration changes the significance of some other consideration (in this case the 1% chance of getting $0) contradicts separability and thus atomism.[74]

By contrast, sometimes what counts as a reason for ϕing in one context may count as a reason against in another (or as no reason at all). Suppose I am organizing a party. I've already invited a number of people, but have some space left. I'm considering inviting Mary and inviting Paul. That Mary is lovely at parties is a reason for inviting her. Mary's presence counts as reason in favor of the party going well. Likewise, that Paul is lovely at parties is a reason for inviting Paul. Paul's presence counts as a reason in favor of the party going well. But suppose that Mary and Paul hate each other, such that their combined presence would be disastrous. Then, given that I've invited Mary, Paul's presence would no longer count in favor of the party going well, but would count in favor of the party *not* going well (and vice-versa). If I want the party to go well, then I have reason to invite each of them, but their combined presence doesn't present two reasons for thinking the party will go well (it does the exact opposite). That is because these are considerations whose significance is open to revision given further considerations, change of context, etc. It's worth noting a distinctive difference here between, on the one hand, a reason being overruled by a stronger reason: if telling the truth would cause significantly more harm, that does not mean that the recognition that ϕ is a lie ceases to be a reason against ϕing; rather, we just have a stronger to ϕ than to not. By contrast, in the party case, Paul's presence makes Mary's presence a reason against the party going well (and vice-versa).[75] If one thinks that *all* considerations are of this form, then one is a *holist* concerning considerations. This is exactly to deny the formulation of atomism given above.

I've described atomism as the position that if somethign is a reason in one context, then it is always a reason (in any context whatsoever); and holism as the position that *no*

[74]I will discuss this below, but decision theorists often make a move reminiscent of minimalists. Namely, they will try to redescribe the situation so that the outcomes in L_1 and in L_3 *are different*. That is, we must understand the outcome of receiving $0 in L_1 as being somehow *worse* since it's the only outcome where we could get $0. Other strategies involve combining states in unintuitive ways. At any rate, this is not a dissertation that is concerned with the nitty-gritty of decision theory. I only make use of it to give an example of an account that explicitly endorses atomism.

[75]Generalists will try to re-describe this sort of case as well. Mary's presence and Paul's presence continue to be reasons in favor of the party going well, but *their combined presence* is a much stronger reason against it going well.

considerations be so. But surely there are intermediate positions? This is indeed the case. For example, it is possible to hold while considerations may gain a different significance in combination (e.g. a reason for becomes a reason against or no reason at all) that this happens in completely predictable ways (i.e. according to principles). This is the view of Horty described above. Horty seeks to describe a view that meets the letter of holism: what is a reason alone may cease to be so with the addition of further reasons, without going in for particularism: that there are no principles. As I explain in the next sub-section, I'll call this sort of position "moderate generalism".[76]

Similarly, one could hold that many (or most) considerations are atomic, but there are still a class (or room for) considerations which are entirely holistic. An example of this kind is Ross' ethical intuitionism. Ross holds that there is a set of core considerations which generate verdicts in a large number of cases, but that there is still some room around the edges where principles cannot determine what we do. Perhaps ϕ is lying but to not-ϕ would be to break a promise. How do we decide whether or not to ϕ in the face of conflicting considerations (that are otherwise decisive)? Ross thinks that around the periphery there are no principles that can guide us, but that we are instead guided by our intuitions concerning ethics. In general, the sorts of things that Ross thinks we appeal to when intuiting verdicts in such cases are contributory reasons. Following Horty,[77] I call this position "moderate particularism". This sort of view is best appreciated (perhaps) in jurisprudence—where, indeed, intuitionism is a view about how judges reach verdicts. According to this view judges have a number of principles that they weigh when making verdicts. These principles are abstracted from case law and perhaps scholarly work on that case law. But there's always some distance between the principles appealed to in case law and the messy details of a particular case before a judge. What bridges this gap (or could be said to bridge it by intuitionists) is some capacity of the judge to intuit how things come down in this case given the particularities (the *dispositive* facts). It's not that the judge lacks reasons for his decision (often times, she has too many, that's the issue!), but that the reasons alone—though some of them may be

[76] I call it this because while it endorses the idea that what is a reason alone might cease to be so in the presence of additional considerations, it does not endorse this universally. That is, on Horty's view, there is a way of describing a reason such that it is immune to further revision.

[77] Horty (2012, p. 147ff.).

decisive—do not yet yield an overall verdict in the case.

4.2.3 Particularists and Generalists

So far I have given the outline of the debate and rehearsed two important distinctions. Before I turn to arguments for particularism, let me briefly give the layout of the debate. *The* important question is whether the considerations present give overall verdicts in accordance with principles. As discussed above, we can either have principles which tell us directly what overall verdicts are (as in e.g. Kant, Scanlon, etc.) or we can have principles that tell us how various contributory reasons combine.

Particularism consists of two commitments:

1. The priority of contributory reasons to overall reason (that the latter must be understood in terms of the former)
2. Holism concerning reasons (in particular contributory reasons)

While there are several proponents of particularism, I will focus exclusively on the version offered by Jonathan Dancy.[78]

Generalism is the negation of this view, but I'll use the label to refer to a specific version of that, namely one that is committed to:

1. The priority of overall reason to contributory reasons (no account of the latter is needed for ethics)
2. Atomism concerning reason

Of the views I've sketched so far Kantian deontology is the clearest example of generalism, but Scanlon also presents such an account.

Horty, as explained above, agrees with the particularist that contributory reasons should be given priority, but disagrees concerning holism (as I've cashed out holism). I'll describe Horty (or the default logician) as *not quite endorsing atomism*, and call such a position (or accounts like this) "moderate generalists" since they agree (with the generalist) that where reasons are given to support what we ought to do (i.e. overall reason: verdicts), they must

[78]Dancy (2004b). For another example, see Lance and Little (2006), though according to my classification, I think we should regard Lance & Little as "moderate particularists".

involve the use of principles. In particular, Horty thinks that principles tell us which of the contributory reasons we have on particular occasions count as decisive. On the basis of such principles, we arrive at overall verdicts. The decision theorist operates similarly (though she needn't isolate some single reason, rather it is the weighing of contributory reasons that provides a verdict).

Ross' intuitionism also agrees with the generalist that overall reason is to be given priority. But there is some degree of disagreement concerning atomism. Such a view thinks that while there are moral principles and these can decide the issue in a great number of cases, there will always be room between what principles guide us to do and the particularities of a particular occasion. For this reason, I describe Ross (or the ethical intuitionist) as *not quite endorsing holism*, and call such a position (or accounts like this) "moderate particularists" since they agree (with the particularist) that sometimes we cannot fully account for what we have overall reason to do via principles.

It's worth remarking on a way in which the generalist and the particularist are rather close together. It might seem at first glance that Kant and Dancy, for example, could not be further apart. Kant thinks that relatively simple principles concerning features intrinsic to actions yield normative verdicts concerning those actions. Dancy thinks that reasons intrinsic and extrinsic to an act may combine in unaccountably complex ways to yield verdicts on particular occasions. But Kant thinks that the categorical imperative can only yield a very small set of ethical prescripts: namely the "strong oughts", i.e. verdicts concerning what we are prohibited, required, or permitted to do. But these principles do not tell us any "weak oughts", that is, what we have reason to do, but are not required to do. Generalists of this type agree with the particularist that how to resolve "weak oughts"[79] is not something that we can give a precise account of. The disagreement concerns whether there is such a thing as a "strong ought", i.e. a consideration which, on all occasions, issues a decisive reason (and thus an overall verdict). Kant believes these to be important for constraining and providing guidance for weak oughts. For example, we are never required to perform any particular act of charity (we are permitted, in each case, to abstain from charity) but are required to perform *some* act of charity (we are prohibited from abstaining in every case). In a certain

[79]In Kant's saying: how to resolve imperfect duties.

sense, the purest form of generalism and the purest form of particularism are not terribly far apart. They disagree about a rather small issue: whether there is such a thing as a strong ought, i.e. the existence of reasons which are decisive on all occasions.[80]

4.2.4 The Holism of the Contributory

With the broad outlines of the debate on the table, I'll now look at some particular arguments. The central argument in favor of generalism is simply that the idea of a reason presupposes some notion of generality. When we explain why, on a particular occasion, we ϕ (or for what reason we ought to ϕ), we do so by appealing to what we had reason to do. In virtue of appealing to such reason, the thought goes, we invoke some property that our action has and that property that explains the reason relation (and the relevant principle). Even if we don't think that moral principles take us all the way to a fully determinate course of action in every situation, they surely provide a great deal of constraints concerning what we may or may not do. The simplest version of generalism—here my example is Kantian deontology—holds that intrinsic features of an act (or more accurately: rules for acting) determine overall verdicts directly. More nuanced generalists (example here: Scanlon) hold roughly the same, but allow that these features—or at least the principles that refer to them—can be extremely complex.

Against this, the particularist argues, when we explain what we had reason to do on a particular occasion, we are saying what we had *overall reason* to do. Next, either reason is provided at this level (i.e., without recourse to the contributory, as in the accounts surveyed above from Scanlon and Kant), in which case, the argument is simply that these accounts get the wrong verdicts. Or, what we had *overall reason* to do must be explained in terms of some further reasons, i.e. contributory reasons. Next, the particularist argues, that because of the holism of the contributory, no account can be given (no principles can be stated) that explain why these reasons combine on this occasion to yield this overall verdict: except the judgment that they in fact do on such an occasion. Any account that purports to do so

[80]This paragraph is meant to mirror a paragraph from the first half of the chapter in which I described how the radical contextualist and the semantic minimalist agree about the complexity of pragmatic phenomena, but disagree about a rather small issue: whether there is a small island of semantic content form which we can gain some purchase.

either (i) gets the wrong verdicts, or (ii) misdescribes what counts as a reason for what. Central to all of this is the argument that contributory reasons are holistic.

The argument that contributory reasons are holistic begins by distinguishing different sorts of considerations that may combine at this level as well as how such considerations interact with each other. This list shouldn't be taken as exhaustive (in fact it is easy to hypothesize many more such considerations and some of my work above shows how we can easily find such considerations—e.g., enabler disablers, attenuator enablers, etc.):

Favorer/Disfavorer: These are considerations which count in favor (or against) something (for example an action or conclusion to be drawn). Typically, these are contributory reasons, but their status as such is subject to change.[81] When an action is justified (or not) these are the sorts of considerations that are cited *in favor* (or against) that action. For example, "I promised my friend I would help her" is a favorer (it speaks in favor of helping) and indeed, often this is reason enough to help her (i.e. to conclude that I ought to help her).

But favorers needn't always be reasons. For example, if I promised my friend I would help her, but a new episode of my favorite TV show is premiering, the latter speaks in favor of not helping. Nevertheless, we think in light of the former it ceases to be a *reason* against helping. It might explain some resistance I have to overcome to helping—or, if the help seems completely unneeded, the basis for some resentment—but it is not a reason against helping.

Enabler/Disabler: These are considerations which enable (or disable) another consideration's role as a favorer (or disfavorer). For example "My promise was made under duress" is a disabler of the favorer "I promised my friend I would help her". It is a *disabler* (instead of a disfavorer) because it renders the original favorer *invalid* (that I promised no longer speaks in favor since the promise was made under duress). Likewise, that the promise was made under duress does not speak *against* helping her. Even though my promise is no longer a reason, she is still my friend and that might be reason enough to help her.

[81] One significant difference between a "favorer" and a "reason" is that something can speak in favor even if it ceases to be a reason.

Intensifier/Attenuator: Finally, there are considerations which serve to strengthen (or weaken) other considerations (in particular favorers/enablers). For example, the considerations that "My friend will experience only a slight inconvenience if not helped" and "My friend has a number of other people around her capable of helping" do not count *against* helping her. I cannot actually offer these as reasons *against helping*.[82] None of these count against helping, instead they weaken the reason I already had: that I promised her. But if I have a competing reason—a reason which, thanks to these attenuators, is now the stronger reason—then *that* reason (disfavorer) is the reason for not helping.

For example, taking an opportunity to get ahead on grading, we might think is not a strong enough reason to break a promise to my friend. But if the force of the promise is weak enough, it may be. Often our intuitions are that it is *no reason at all* (except when it is, i.e., when the force of the promise is weak enough). An explanation of this is that moral reasons often silence other reasons (though they don't silence the favoring: it remains true that the opportunity to grade speaks in favor of not helping, even if—according to this intuition—it is not a reason).

A consideration could be any of these or even multiple at once. For example, suppose that Mary is not only great at parties, but also gets along great with Sam. Ordinarily Sam's presence at a party doesn't speak in favor or against the party going well (he's a pleasant enough person). However, Mary and Sam get along great and her presence is not only great on its own, but also makes Sam's presence count in favor of a party going well. Here, Mary's presence at the party is not only a favorer (of the party going well) but also an enabler of the status of Sam's presence as a favorer (without Mary, Sam's presence was not a consideration that spoke in favor of the party going well).[83]

What this list is meant to preclude is **two sorts** of generalist strategies for accounting for considerations. Concerning generalists who deny accounting for the overall in terms of

[82]Though note that we do invoke these when challenged. Why? Because we cannot contest the promise made, but we can contest its strength as a reason.

[83]It's worth noting here that the way various considerations interact precisely rules out a criterion like separability. Mary's presence and Sam's presence are not separable as far as decision theory is concerned. A possible work around is to divide the states into four (Sam alone, Mary alone, both, neither), but this might seem like an artificial fix, and even more so in more complicated cases.

the contributory (e.g. Kant, Scanlon), the argument is meant to show that the considerations which may be relevant for determining the moral verdict of an action are potentially inexhaustible. Thus, determining whether an action is right/wrong simply by examining the action—and even if such a principle has a number of unless/except/etc. clauses—is not possible.[84] If this is right, then no finite set of principles at the overall level could be formulated without those principles containing at least some false verdicts (since we could always add more considerations to generate counter-examples).

For example, the sort of principles that a contractualist like Scanlon would allow might be quite complicated. While he thinks that we can talk perfectly sensibly of quite general principles like "thou shalt not kill" or that one ought to keep one's promises, we should understand these as a kind of crude approximation of what our principles actually say. He writes:

> "even the most familiar moral principles are not rules which can be easily applied without appeals to judgment. Their succinct verbal formulations turn out on closer examination to be mere labels for much more complex ideas. Moral principles are in this respect much like some legal ones. The constitutional formula "Congress shall make no law abridging freedom of speech, or of the press" may sound like a simple prohibition. But the underlying idea is much more complicated.
> [...]
> In making particular judgments of right and wrong we are drawing on this complex understanding, rather than applying a statable rule, and this understanding enables us to arrive at conclusions about new and difficult cases, which no rule would cover.
> [...]
> When we judge a person to have acted in a way that was morally wrong, we take her or him to have acted on a reason that is morally disallowed, or to have given a reason more weight than is morally permitted, or to have failed to see the relevance or weight of some countervailing reason which, morally, must take precedence. Each of these judgments involves a principle in the broad sense in which I am using that term. [...] But we make such judgments by drawing on our understanding of why there should be a moral constraint on actions of the kind in question (why principles that left us free to do as we liked in such situations are "reasonably rejectable") and of the structure that that constraint takes (in what way we can be asked to take the relevant interests into account). When, in the light of our best understanding of this moral rationale, we make a judgment about the sufficiency of the reasons for an action in a particular case, this judgment is guided by, and expresses, our understanding of a moral principle." (Scanlon, 1998, p. 199–201)

[84] Note that this claim is not equivalent to a denial of atomism. This claim is, however, a denial that we can do away with the contributory or not treat it as basic. A consequentialist, for example, can make sense of the idea that there might always be more considerations relevant to verdicts concerning the rightness/wrongness of an action because there are always potentially further relevant consequences of performing the action.

What Scanlon means here when he distinguishes a “statable rule” from a complex understanding of a moral principle is that moral principles needn't be the sort of thing that one could give instructions to non-moral creatures to follow. It may be that certain principles may be sensitive to whether an action would be *cruel* or *courageous* and perhaps whether, in general, an action is of this type is not something that can be explained by appeal to further features (e.g. so-called thin concepts). This is fine, particularism doesn't turn on whether there are thick ethical concepts (such as cruel, courageous, etc.).[85] What Scanlon *is* endorsing is a set of principles—perhaps some of them quite complex or involving concepts which require some level of judgment—that explain why an action is right/wrong on a particular occasion. This means that while we might not be able to state such principles with exactness on any particular occasion, we can always explain what made an action right/wrong by appeal to *general features* of that action (and so to features which would have the same significance on other occasions).

Now, as complex as these ideas are, surely they must be finite, in the sense that we could reach a description of an action that is simply decisive. That ϕ is a killing might not be such a principle, but perhaps:

- ϕ is wrong if ϕ is an unjustified killing.
- ϕ is wrong if ϕ is an unjustified killing undertaken with the intent to kill.
- ϕ is wrong if ϕ is an unjustified killing undertaken with the intent to kill for reasons that *could* be reasonably rejected.
- etc.

What this brand of generalism amounts to is the claim that we can always find such a privileged description of an action that simply settles it (come what may) that an action is right/wrong. What particularism amounts to is a rejection of this, that there is always some further consideration which unseats the above description as *decisive*.[86]

[85] Though some take the existence of thick concepts to be decisive of the matter. For example, whenever we ought to ϕ, we might think that the description of the action that settles it is simply: “ϕ is right”. But then the issue seems to turn on whether there is some other description of ϕ that could be invoked as a reason. I am generally avoiding the thick-/thin-concept framing of the debate.

[86] I don't want to stake too much on it, but in case the reader has doubt that we could actually find considerations to unseat the third sort of reason, here is a shot at explaining some further considerations which might alter the significance of the fact that “ϕ is an unjustified killing undertaken with the intent to kill for reasons that *could* be reasonably rejected”.

The second strategy it precludes is any sort of account which seeks to show how various contributory reasons may combine to yield an overall verdict. This is because not only are the potentially relevant considerations, potentially inexhaustible, but further considerations do not simply consist of further reasons, but may in fact change what previously counted as a reason or the way in which reasons interact with one another. As I explained earlier, this is straightforwardly rejection of separability—we cannot evaluate reasons independently of what other reasons, outcomes, or considerations are present—but this sort of objection, also poses problems for the other view I considered which places contributory reasons prior to overall reason: default logic. To put the matter explicitly: the charge of the particularist is that either (i) the generalist gets the wrong verdict or (ii) is forced to misdescribe what is a reason for what (at the contributory level).

Let us start with the rational choice theorist and Allois' paradox referenced above. The particularist charges that either (i) the rational choice theorist gets the wrong verdict (that L_2 is not to be preferred to L_1) or (ii) the rational choice theorist must misdescribe what counts as a reason for what. One such strategy is to redescribe the \$0 outcome in L_1 such that it's not merely understood as \$0 but \$0 *plus being the only zero-money outcome*.[87] There are several problems with this suggestion. First, features of the entire situation aren't typically invoked as *reasons* for or against an action. To make this explicit, consider the following justification in support of ϕing: A provided reason to ϕ and I had no stronger reason against ϕing. *That I had no stronger reason* is itself not a reason to ϕ, it is a statement that follows the overall verdict concerning what to do. That, on the whole, I had overall reason to ϕ

In general, that ϕ is illegal is a reason against doing ϕ: this is true both in general and for Scanlon. Since, all things considered, following the law is a principle that could not be reasonably rejected (ignoring small violations; e.g. jaywalking). But if a law is particularly immoral, then that ϕ is illegal might be a reason to do ϕ. In fact that ϕ might be righteous precisely because it is the unjustified killing of someone motivated by some structural deficit in the law. Some of the murders carried out by the abolitionist John Brown, or (were it successful) the attempted assassination of Henry Clay Frick, could be defended as historical examples of this. That the slavers or Frick were acting perfectly legally (which is perhaps to say according to principles that could not be reasonably rejected—though, the question of "by whom" looms large) is precisely the reason for the political violence.

[87] Compare some of my remarks in a previous chapter about treating structural features of reasoning as particular considerations—I mean my decision theory examples in Ch. 2, §4.

NB: I am treating outcomes as reasons for actions for decision theorist. But the decision theorist might simply abandon the idea of reasons for action. What provides *overall reason* is something about how all of the considerations in a situation combine. But we can't single any considerations out as reasons for/against taking any particular action.

is likewise not a responsive answer to the question, "why did you ϕ?" Second, it seems to misdescribe our understanding of the situation. The \$0 outcomes (on a roll of 1) in L_1 and L_3 are understood as the same outcome: \$0, but because of the additional considerations present (namely the assurance of \$1m in L_2 or the likelihood of \$0 in $L_{3/4}$), the weight of this reason changes.

We may make a similar criticism concerning default logic. Let me put a few more details on the table. Officially, X is a reason for Y (on a particular occasion) iff there is a default rule of the form $X \mapsto Y$:

> "Our discussion throughout this book will therefore be based on an analysis according to which *reasons are identified with the premises of triggered defaults*; and we will speak of these triggered defaults, not as reasons themselves, but as *providing* certain propositions—their premises—as reasons for their conclusions." (Horty, 2012, p. 27)

Reasons are identified with the premises of default rules. This is just to say that what it means for there to be a default rule is for there to be a reason on a particular occasion. That ϕ is a case of lying is a reason against ϕing. This is *identified* with the default rule:

$$L(\phi) \mapsto O(\neg\phi).$$

Where $L(\phi)$ means to ϕ is to lie and $O(\neg\phi)$ is to be read as one ought not ϕ. $L(\phi)$ is a consideration that counts against ϕing. The most straightforward objection against default logic is that it is unable to account for the fact that on one occasion a consideration may be weightier than another—that one do that which is likely to cause substantially more good than harm is weightier than that ϕ is a case of lying—may be otherwise on another occasion or with the addition of some further consider: for example, a prior promise not to deceive concerning the specific issue at hand. So, often in such cases (as with cases of agglomeration) default logic seems to get things wrong. Default logic does allow for two work-arounds, but both misdescribe what counts as a reason for what.

For example, that ϕ was promised ($Pr(\phi)$) is a reason to ϕ (one ought ϕ: $O(\phi)$):

$$d_1 : Pr(\phi) \mapsto O(\phi).$$

That ϕ was made under duress ($D(\phi)$), however, is represented by the default logician as a reason against ϕ:

$$d_2 : D(\phi) \mapsto O(\neg\phi),$$

and then d_2 is given a higher priority than d_1. But this is wrong on two counts: first, that ϕ was made under duress is *not a reason* against ϕ.[88] It simply disables what was a reason (namely that ϕ was promised). Second, though d_2 is ranked higher, d_1 is still triggered, so still counts as a reason to ϕ, even if, on the whole—i.e. at the overall level—we have reason not to ϕ. This is a problem since it doesn't properly distinguish a case where a reason is disabled (as it is here) from cases where that reason is instead weakened or where a competing reason is simply stronger (the last of these seems most easily accounted for by the default logician).

4.2.5 The Common Presupposition

Among those who think that moral principles can take us all the way to moral verdicts, there is a divide between those who think that the only principles needed concern overall reason (i.e. concern those moral verdicts directly) and those who think that contributory reasons play some role in explaining those verdicts.. These accounts seem unable to account for complex ways that additional considerations may alter these verdicts. What they are committed to is the idea that there are descriptions of action which ground overall verdicts concerning what is right/wrong (i.e. what we ought to do). These descriptions could be quite complex—as Scanlon explained, "thou shalt not kill" might be "shorthand" for a much more complex idea in the same way that "congress shall make no law [...] abridging the freedom of speech" is shorthand for a much more complex idea—but there are nevertheless descriptions that, when present, give definitive verdicts concerning right/wrong. Because of the way that additional considerations may influence these verdicts, it is worth considering the form these principles must take (even if quite complex). "thou shalt not kill" can't work because there are cases where killing is permitted. So any principle prohibiting killing would need to acknowledge cases where killing is permitted as well as exceptions to those principles and so

[88] This is because of how "oughts" are cashed out by Horty. We may avoid this first problem by reformulating the rule as $D(\phi) \mapsto \neg O(\phi)$, but the second problem still looms.

on. While these principles are not formulated as such, it is not difficult to see that what this sort of generalist is after is a set of individually necessary and jointly sufficient conditions for an action to be right/wrong.[89] This is because a principle which prohibits an act (says that it is wrong) and says of all and only the wrong acts of this type that they are wrong is equivalent to a list of individually necessary and jointly sufficient conditions. Since the generalist thinks that this is what is required of a principle to reach verdicts concerning right/wrong and the particularist thinks that this cannot hold, we can characterize their shared ground via the following conditional:[90]

If an act is governed by principles, then the recipe for ascertaining those principles must be contained within the act itself.

Next, let us consider Ross' account. Recall that on Ross' account moral principles operate at the level of overall reason. *Prima facie* reasons identify features of actions which determine immediately whether that action is prohibited/mandatory/permitted, and thus issue normative verdicts. What distinguishes Ross from Scanlon and Kant is that Ross understands these principles to be *prima facie* or *pro tanto*. That is, that they hold, all things considered, or so long as no other principles conflict with our verdict.[91] So *prima facie* features of an action tell us normative verdicts directly, *except* where such features may come into conflict. In such cases, features of a situation may need to help us decide. These features are contributory reasons. Here's a famous example: I am contemplating whether to switch trolley tracks and kill one person instead of five. My action is pulling the lever, two principles seem to conflict:

- To pull the lever would be to kill, so *prima facie* I must not pull the lever
- To pull the lever would save the most lives, so *prima facie* I must pull the lever

[89] Or perhaps action-type. I'm allowing, as I think we ought, that a generalist can tolerate moral conflict. That is, an act might be both required and prohibited under two descriptions (even if we had the kind of privileged descriptions the principles demand).

[90] Recall that we said that an act has moral content means that there are moral principles concerning that act, i.e. we can say things concerning the goodness/badness or rightness/wrongness of the act in abstraction from particular occasions where the act may be undertaken.

[91] In a long footnote in a previous chapter, I credited Chisholm (1964, 1966) with introducing the language of "defeat" into philosophy. Chisholm's own work here was inspired by Ross, but my motivation in observing this here is to remark that it would not be inappropriate to call the phenomenon in question, "defeat", except that Ross didn't have such vocabulary available to him.

In this case, some features of the situation must tell us how to resolve this conflict. We must determine whether this is a case where non-maleficence overrules our duty to promote a maximum of aggregate good, or vice-versa. Such information is not strictly contained within the act itself, but must be intuited from our previous grasp of such principles. Even if this further information cannot be captured in principles—and for this reason I counted Ross as a "moderate particularist"—the act itself still provides us with a recipe for ascertaining that information. Where *prima facie* duties conflict, the particular situation resolves the conflict in favor of one duty or another. At any rate, it is clear that Ross is also holds this common presupposition.

Next, let us consider those who think that overall reason must be explained in terms of contributory reasons. This includes the particularist as well as moderate generalists such as decision theorists and default logicians. As explained above, the particularist may lodge two complaints against moderate generalists: first, they rely on principles for combining contributory reasons to reach an overall verdicts (and these principles are sensitive to the same sort of particularist' concern as other principles are); second, they require principles for specifying what counts as a reason for/against an action. I will focus on the latter of these to find the shared ground. My strategy is to show that both the decision theorist and default logician conceive the act itself as providing the recipe for determining a normative outcome.

Let us start with the example of the decision theorist. Recall that outcomes provide reasons for acts.[92] The strength of these reasons is perhaps weighted to their probabilities and some specific criterion is used to determine what we have overall reason to do. There are two closely linked assumptions that the decision theorist relies on which will get us quickly to the shared presupposition. The first is that acts are understood as functions from states of the world to outcomes.[93] This means that to have an understanding of an act is to have an understanding of all the reasons for or against undertaking that act. Further, since the outcomes are separable, no further considerations can change the status of any outcome as

[92]Many prefer to call these "prospects" or "options", but I am being somewhat loose where possible.

[93]Savage (1954) defines them this way. Broome (1991) understands this requirement to follow from separability. A slightly more nuanced argument is required for Jeffrey (1990), but I think that what I am pushing in this section applies equally well there.

a reason for or against the act.

The second assumption is what is called the "rectangular field assumption".[94] This assumes that the space of acts is the space of functions from states of the world to outcomes. Given further assumptions about the preference ordering on acts—some of which are taken to be guaranteed by separability—this gives us an immediate answer concerning what act we ought to do (namely the best one available).[95] Of course, the particularist will deny both of these assumptions. We cannot glean from an act what reasons count in favor or against it (further considerations can always change those). Likewise, we cannot say whether one act is to be preferred to another for the exact same reason. In one context, A may be preferable to B, but given further considerations the opposite may hold.[96] Then again, it seems clear that the common ground is the same:

If an act is governed by principles, then the recipe for ascertaining those principles must be contained within the act itself.

This is because *what an act is* is nothing other a function from states of the world to outcomes (i.e. reasons for/against the act).

Finally, consider the default logician. Recall that above I argued that while Horty's default logic is able to accommodate some of the data of the particularist, it only does so in an artificial manner. The problem is that a considerations status as a reason cannot be infirmed by further information. Even moving the goalposts—so that the status as a *good* reason is what is at issue—this change is similarly limited. Regardless, is it possible to construct a recipe for ascertaining normative verdicts from the action itself? I think so. Given how Horty's theory is constructed, a default can only be triggered if the premise is contained within (or logically implied by) the set of initial considerations (i.e. if the premise of a default would only be triggered by some chain of "default reasoning", then it is not triggered at all).[97] What this means is that given a scenario and an action ϕ, we can list

[94] Broome (1991).

[95] To guarantee a normative verdict, we need to require transitivity and completeness of the preference ordering. But completeness is potentially controversial (even more so than transitivity).

[96] Allais' paradox is an example of this. See above.

[97] This is because classical—fully structural—logic is the engine that powers default logic. If we allowed chains of default reasoning to determine which defaults were triggered defaults, then any conflict among default rules would generate a contradiction and hence explosion. In fact, in order to get the right verdicts,

exactly which defaults have ϕ as a conclusion (and are triggered) as well as which are in conflict. This means we can simply write a rule that specifies immediately the necessary and sufficient conditions for ϕ.[98] Since we needn't worry about chains of default reasoning, it should be clear that the act itself (or those defaults with which it is involved) provide us a recipe for ascertaining its moral content. That is, the default logician is likewise committed to:

If an act is governed by principles, then the recipe for ascertaining those principles must be contained within the act itself.

4.2.6 Substructural Reasons

Next, let's fill in some of the details of this common presupposition. It seems that all participants agreed that if an act is to be governed by principles, then the act itself needs to inform us how to ascertain those principles. That is, either the description of the act itself supplies an immediate answer to the question of whether or not the act ought to be done (as was the case for generalism), or the act tells us what sorts of additional information would be needed to reach such a conclusion (as was the case for moderate particularists). As in the first half of this chapter, I'm going to move quickly to fill in some of the details here. I can't claim to offer an argument that the participants in the debate *must* conceive of things as I am about to claim they do—after all, I am about to show a different way to conceive of things—but it seems to be a pervasive way of framing the debate that I don't think is terribly controversial.

The first move is to fill in the details of "moral principles". All of the participants in the debate take it for granted that moral principles single out features (principally: of an

an awful lot of pre-arranging is required—compare the role of "speaker intention" for the minimalist—that is, we have to carefully select which rules and which propositions go into a particular scenario. It can make a substantive difference whether we describe a scenario as containing $\{P, P \supset B\}$ or $\{P, P \supset B, B\}$ (even though their logical closures are identical). In both cases defaults with B as a premise will be triggered, but only in the latter case, will the conclusions of such defaults trigger the premises of further default rules (e.g. if $d_1 : B \mapsto C$ and $d_2 : C \mapsto E$, then both trigger d_1, but only the latter triggers d_2. But, as I've said, the goal of this chapter isn't a knock-down refutation of Horty. My increased attention to his account is proportional to its sophistication and compellingness.

[98] Simply start with the triggered defaults. We may ignore anything *lower* (=less preferred) on the preference ordering. Then we simply go up the preference ordering looking for defaults which conflict.

act) as reasons for or against a particular act.[99] For example: "to ϕ is to lie" is a feature of ϕ. A principle such as "Because "to ϕ is to lie", one ought not ϕ" singles out this feature as a *reason* to not ϕ. Typically principles are bridge principles between normative and non-normative vocabularies (though there can also be principles linking purely normative vocabulary, for example). In the simplest case (as just referenced) whether that ϕ is a lie may be stated without use of normatively laden terms like right/wrong and so on. For example, in law judges invokes principles which connect matters of fact with matters of law: if these facts obtain, then these legal consequences follow.

A recipe for ascertaining that content means some way of moving from acts to normative verdicts concerning those acts. This means that we need some finite set of principles for each act and some finite procedure for adjudicating these principles (i.e. some procedure for conflict resolution or weighing). In the simplest case (what the generalist endorses), there is a privileged principle for each act that delivers an immediate normative verdict. A slightly more complicated procedure might have it that an act has attached to it a large number of reasons (for and against that act) and that the weighing of these reasons provides a normative reason concerning that act. Horty provides a procedure whereby some contributory reason is deemed decisive, and thereby provides a verdict at the overall level. What is important, however, is that there be some finite set of principles governing each act and that these principles—or whatever lacunae they generate—are entirely predictable from the nature of the act itself. This need not entail that a rational agent could list them—compare Scanlon's remarks concerning "simplified" principles such as "thou shalt not kill" with the complex principles of which they are simplification—just that the sort of information relevant to the normative evaluation of an act is predictable from the nature of the act itself.

The second move is that because there is a finite list of principles governing normative verdicts for each act, we can generate necessary and sufficient conditions for normative verdicts. This move might seem too fast; in fact, perhaps it seems even faster than when I

[99]This is not strictly speaking true! Decision theorists have an opening for denying this. They provide principles for normative verdicts that needn't be thought of as selecting certain considerations are "reasons" for/against any particular act. They can instead simply claim that it is the way that all of the considerations hang together that determines the outcome. Though "the way they hang together" is precisely, mathematically describable. This way of cashing out decision theory places the decision theorist much, much closer to my view, though they still ultimately endorse a fully structural view of what principles must look like. Broome seems to understand decision theory as weighing various reasons provided by outcomes.

invoked a similar move in the first half of the chapter. But the idea is this: for each act, there is some set of principles and some set of further principles for adjudicating those principles. This gives us a set of considerations (perhaps some of which must be formulated disjunctively or conjunctively; such considerations may also involve the use of so-called "thick concepts") that are individually necessary and jointly sufficient for a normative verdict for a particular act. For the decision theorist, each act comes with a set of reasons for/against it, and then further principles tell us how to use these to generate verdicts.

As before then, we can represent this as follows (where $\gamma_1, \ldots, \gamma_n$ are those considerations and ϕ is the action to be performed):

$$\begin{aligned} \gamma_1, \ldots, \gamma_n &\vdash \phi \\ \phi &\vdash \gamma_1 \\ &\vdots \\ \phi &\vdash \gamma_n. \end{aligned}$$

Again (for simplicity), I write $\Gamma \vdash \phi$ and $\phi \vdash \gamma_i$ for $1 \leq i \leq n$. These exhaust the principles governing normative verdicts for ϕ. These considerations are clearly structural: they are clearly monotonic and transitive.[100]

One feature here worth remarking on, of course is that we may require different necessary and sufficient conditions for the verdict that ϕ ought to be done as well as the verdict that ϕ ought *not be done*. Since, supposing it's not the case that ϕ ought to be done whenever the above conditions *don't* hold, it does not follow that ϕ ought not to be done.[101]

Regardless, these are the two features that the particularist wants to deny. It may be that $\Gamma \vdash \phi$, but a further reason could defeat this. It may be that $A, \Gamma \nvdash \phi$. Likewise, it may

[100] I again leave to the side contraction and reflexivity for the same sorts of reasons. Though I will note *here* that denials of contraction actually seem to be much more plausible on the practical side. Since two instances of a reason may agglomerate and provide more reason than one instance (i.e. provide normative verdicts that one instance would not provide). But I've been avoiding any kind of metaphysical debate throughout this chapter. In particular, if ϕ would save two lives, does that provide two reasons? Each person has their own claim? Or do the two claims combine into their own reason (a kind of joint venture)?

[101] A strong analogy exists with the first half of the chapter here of course. Since we could read consequence there normatively. Suppose that '$\vdash$' preserves truth and that we understand assertion normatively (e.g. in terms of a truth or knowledge norm). Then we could simply read $\Gamma \vdash A$ as A may be asserted (provided Γ). But, there is a clear difference between when we may assert $\neg A$ and when we may not assert A (the former requires sufficient information; the latter not).

be that this further feature is *no reason at all*. Hence on one occasion we may have $\Gamma \vdash \phi$, but on another $\Gamma \not\vdash \phi$ (without the addition of any reasons). *A quick aside:* generalists (and the two moderate camps) usually simply count A or A together with some consideration) as a kind of reason to get around this. But this is precisely what the particularist means to deny.

Likewise transitivity. Suppose that Γ provides reason to ϕ. Suppose further that ϕ requires (or provides reason) to A. Does it follow that Γ provides reason to ϕ? Let's look at a concrete example. I have promised my friend I will pick her up from the airport ($Pr(A)$). So my promise provides reason for the normative verdict that I ought to pick her up ($O(A)$):

$$Pr(A) \vdash O(A).$$

But, my car has insufficient gas (E). So I *must* put gas in my car in order to pick her up from the airport ($O(G)$). Thus:

$$O(A), E \vdash O(G).$$

Does it follow that:

$$Pr(A), E \vdash O(G).$$

The particularist will deny this. Promising my friend doesn't provide a reason to put gas in the car. The reason I must put gas in the car is because I must pick her up and my car requires gas. That's not to say that there's no story to be told here. That I promised her, clearly plays some role here. It generates the obligation to pick her up. But it does not generate the obligation to put gas in the car.

The view I have been advancing in the dissertation, however, allows us to understand how both of these patterns may obtain. If we understand the force of a reason in terms of the normative verdicts it contributes to, then it may stand in such relations that are substructural: non-monotonic and non-transitive. The upshot is that we get an account of how a consideration can be governed by principles without requiring the specification of particular reasons. That is, we understand the principles surrounding a particular normative verdict in terms of what considerations imply it, and together with which considerations, what further obligations or considerations it may imply. Such principles must are fully substructural. We understand the "moral content" of a consideration A (and thus how it is governed by principles) in terms of its contribution to normative verdicts.

It might seem that this requires abandoning talk of reasons when discussing normative verdicts—I return to this objection below—in the same way that we needed to abandon truth conditions in the first half of the chapter. I think here a major dissimilarity emerges. We understand the normative verdicts as providing overall reason, so that the verdict that one must ϕ is based upon those considerations which contribute to that particular verdict. But this doesn't require that we specify which of those considerations are reasons *for that verdict* (in the same way that previously we did not specify what made a sentence true—which considerations determined or constituted truth conditions). Instead, we should understand reasons as that with which we reason. The considerations which support a normative verdict on a particular occasion *are reasons*. Why? Because they are used to reason to the conclusion.

To sum up: we do need to abandon a particular notion of reason: that of a contributory reason. We cannot in general say whether a particular consideration will speak in favor (or against) a particular action. But that doesn't mean that that consideration isn't a *reason*. In virtue of the fact that it is relevant to normative verdicts concerning that act, we must use it in our reasoning. But that we use it in our reasoning is all that we need to call such a consideration a *reason*.[102]

4.3 The Significance of a Consideration

In this chapter I examined two debates in disparate regions of philosophy. My intention was to (i) show a deep similarity in the layout of logical space in each of these debates and (ii) show how my account of substructural content opens up new logical space in both debates.

With regard to (ii), I argued that in both debates, there was a common presupposition:

[102] Titelbaum (2019) notices something similar. In that piece he argues that we may have "reasons" without those reasons being for/against a particular act/outcome. That is to say, we can have reasons without reasons for. He doesn't put it exactly the way I put it, but I think he is pursuing a similar line of thought: we can call them reasons because we reason with them. This locution was suggested by Bob Brandom and helped untangle a number of things in this chapter.

that the only way for something to be contentful, was if the thing in question provided a recipe for ascertaining that content. In both cases this amounted to necessary and sufficient conditions for a privileged status of the thing in question. But this in turn seemed to require a fully structural understanding of that content: that is, a view of content in which that content stands in relations of implication which are monotonic and transitive. Against this, it was argued that the phenomena in question behaves in ways that are clearly *not structural*. If we understand this content in terms of its contribution to good implication (or normative verdicts), then we have a way of understanding how something—i.e. a sentence or action—may be contentful but still stand in substructural relations of implication.

I'd like to close the paper with two thoughts. First, (i, above) I'd like to explain the deep similarity in the layout of logical space of the two debates. I'll end by considering an objection to my contribution to both debates. The objection I hope sheds further light both on the similarities between the debates as well as what it means to go in for a notion of substructural content.

Recall, my claim at the beginning concerning how various features in the two debates lined up (reprinted with labels in Table 12). In both debates, we have what we might call in-

Table 12: The Components of the Two Debates (with labels)

	Radical Contextualism & Semantic Minimalism	Moral Particularism & Moral Generalism
(α)	Sentence	Action
(β)	Truth/falsity of an assertion	(Overall) normative verdict
(δ)	Truth-Conditions	Reason(s)
(ϵ)	Semantic Content	Moral Principles
(ζ)	Contextual Information	Contributory Reason
(η)	Radical Contextualism	Moral Particularism
(θ)	Semantic Minimalism	Moral Generalism

trinsic features. (α) In the first half of the chapter, these were features of a sentence (e.g. its grammar and facts from the lexicon concerning the various expressions in the sentence);[103] in the second half these were features of an action (and not further considerations pertaining to a particular performance of that action). As remarked earlier, often these features are non-normative features of an act. (β, δ) Semantic minimalists held that these features alone determined truth-conditions of sentences and thus the truth or falsity of assertions on particular occasions. Moral generalists held that these features alone determined what counts as a reason for/against normative verdicts concerning that action (reasons for its prohibition/requirement/permission). (δ, ζ) Moderate contextualists agreed that the sentence itself gave us the recipe for determining the truth conditions of a sentence, but that contextual information might nevertheless play a role in determining those truth conditions (though it does so in a predictable manner). Likewise, moderate particularists (e.g. Ross) held that the act itself gave us a recipe for determining normative verdicts, but that we may need to consult further information—contributory reasons—to help adjudicate conflicts generated by features of an act. Though the need for such further information is predictable from the act itself. Moderate generalists (e.g. decision theorists and default logicians) similarly allowed that contributory reasons may play a large role in determining normative verdicts (that either we need principles for weighing these reasons or deciding which among them is to be decisive).[104]

(ϵ, θ) Nevertheless, moderate contextualists and semantic minimalists agree that sentences have semantic content and that the recipe for ascertaining that content is contained within the sentence itself. Likewise, (moderate) generalists and moderate particularists think that there are moral principles governing our actions and that the act itself provides a recipe for ascertaining those principles. (η) Radical contextualists denied this possibility and hence the possibility of semantic content. Moral particularists likewise denied such a possibility and thus the possibility (or at least importance of) moral principles.

What is at issue in both debates is the significance that further considerations have.

[103] We might understand truth conditions as connecting non-semantic facts with semantic verdicts—namely truth/falsity of a sentence.

[104] Is there such a view as a "moderate minimalist" to correspond to my "moderate generalist"? I think so, but I didn't cut at such a fine grain in the first debate. Recanati, for example, allows all expressions to be context sensitive. Thus truth-conditions must be constructed out of contextual information. The analog of the moderate particularist would be views which allow contextual information to play a role in a small set of cases. E.g. those who hold that certain expressions are context sensitive (but not that all expressions are).

In particular, their significance with regard to truth conditions (in the first case) and with regard to reason(s) (in the second case). But, as I've been pushing, we needn't think of the content of a sentence in terms of truth conditions; nor the moral content of an action (i.e. the principles which govern it) in terms of reason(s) for/against its performance. Principles don't need to specify reasons *in this sense*, only the significance of a consideration directly (in terms the contribution it makes to normative verdicts). Thus, we needn't understand the significance of a consideration in these terms either. Instead we may understand the significance of a consideration in terms of its contribution to good implication (or its contribution to normative verdicts). Such a view maintains (with the minimalist and generalist) that there are stable contents of sentences and actions (that there are moral principles!). Nevertheless, it allows that the effect that further considerations have on truth conditions or reason(s) might be radical or unable to be accounted for. This leads into my last point: a lingering objection. The objection is the same for both debates, but I'll state it twice:

Truth is Essential: I've claimed that we can simply abandon truth conditions in favor of content understood as contribution to good implication. But a central idea undergirding the debate is whether we can gain an understanding of the content of a sentence apart from its assertion on a particular occasion, so as to understand what is asserted when it is asserted. But assertion must be understood in terms of truth. Hence, the content we are looking for must be truth-apt (since it is what we get when we assert a sentence). To depart from this is simply unintelligible.

Reasons are Essential: I seem to claim that we can simply abandon talk of reason(s) in favor of principles understood as contribution to normative verdicts. But a central idea undergirding the debate is that principles specify what counts as a reason for/against an action apart from a particular performance. But, that we ought or ought not perform an act on a particular occasion must be rationally grounded in reasons. To depart from this is simply unintelligible.

I contest the centrality of either of these.[105] In any case, I don't think it's impossible to say what makes a sentence true/false on a particular occasion (I've hinted as such throughout).

[105]Assertion and normatively sanctioned performance needn't be understood in these terms.

Even so, it may be that we can't provide an account of truth conditions of a sentence that won't vary radically from context to context. Likewise, I don't think it's impossible to specify reasons (for the same sorts of reasons). But even if we can't say in general what counts as a reason for/against an action in general, that does not tell against being able to say on particular occasions what is playing the role of reason. That is: reasons should be understood in terms of what we reason with. That something is a consideration which contributes towards a rational, normative verdict is all that is needed to say that that consideration is a reason: reasons are what we reason with.

What I'm advancing is a notion of content that is not beholden to truth. Or a notion of principle that is not beholden to reason(s) for/against. [Not the same as denying the existence of either!] But it's worth pausing to note that such a suggestion is radical. I am saying that deeper accounts of these notions might not be possible. The consideration, "to ϕ is to lie", is involved in a number of principles: it contributes to normative verdicts! And I think that for that reason we can call it a reason (we use it in reasoning about whether to ϕ). Nevertheless, we might not be able to say in the abstract where it counts as a *reason* for/against ϕ (though we can on particular occasions do so).

4.4 Conclusions

In previous chapters of the dissertation I put forward a theory of the meaning of a sentence in terms of that sentence's contribution to good implication, where such implication was potentially radically substructural and where this substructurality was understood as *pushed all the way down into the content* of the sentence. In this chapter, I used this theory to expose logical space in two debates. The first of these concerned radical contextualism and semantic minimalism in the philosophy of language. Central to this debate was the question whether a sentence has a literal meaning (and whether this literal meaning is what is said by that sentence on all occasions). I argued that we can understand the diverse implications that a sentence is involved in without requiring a change of content by understanding that content to be the contribution the sentence makes to good implication. This required giving up was a conception of content in terms of truth conditions.

Next, I showed that this notion of content can also be used to help clear up a debate in meta-ethics. There, the question was whether there can be principles which specify reason(s) for/against an action and thereby issue normative verdicts concerning those actions. I argued that if we understand those principles substructurally, then we can find such principles. One thing we had to move away from in so doing, however, was the idea (in general) of a reason for/against an action. As I suggested above, this doesn't meant that we can't talk of reasons: reasons are simply considerations that we reason with. But we cannot in general specify whether a consideration will count in favor or against an action, and so principles do not specify whether individual considerations count in general for/against an action.

Finally, I showed that these debates are analogous in a number of ways and that there is an abstract way of describing each of them that shows this commonality. At the outset of the paper, I suggested that a crude formulation of this is that radical contextualism is just particularism for truth conditions, or that moral particularism is just radical contextualism for normative judgments. I think there is something to this crude formulation, but in the close of the previous section, I showed how a more nuanced understanding gains us some additional insight into both debates. In the radical contextualism debate, we found that a sentence may have a stable content even if it fails to have stable truth conditions. In the particularism debate, we found that an action may be governed by reasons, even if we cannot in general say which reasons count for/against the action (except as they do on particular occasions).

5.0 Conclusion

Let me briefly summarize the results and implications of the results developed in this dissertation. I put forward an account of the meaning of a sentence in terms of its contribution to good implication, where such implication was understood to be (potentially radically) substructural. I showed that despite the lack of constraint on content, I was still able to develop a completely tractable formal semantics for which the usual sentential logical connectives could be introduced. The tractability was further demonstrated by the introduction of simple and familiar proof rules (in the form of a sequent calculus) that were proven sound and complete. I next used this proof system to prove some results of particular interest to inferentialists. Namely, I showed that I can use this machinery to precisify the thesis of logical expressivism. Logical expressivism is the thesis that we should understand logical connectives in terms of their expressive role: that is in terms of their ability to express good implication. I showed that despite the radical substructurality of the consequence relation generated, I was able to develop connectives that expressed *exactly* what followed from what. Thus vindicating not only a goal of logical expressivists (a goal which I made more precise), but also the tractability of the entire project.

In addition to demonstrating that we *can* work with such a formal semantics, I gave some motivation for why we might want to. I did this by focusing on one structural rule: monotonicity. I argued that attempts to account for non-monotonicity via defeasible reasoning—i.e. to account for the phenomenon where by $\Gamma \vdash p$ but it may be that $\Delta \nvdash p$ (with $\Delta \supsetneq \Gamma$)—will not work. They will not work because we are unable to gain a grasp of what a "defeater" is. The best we can say is that additional considerations may change the way that a set of considerations hangs together such that an implication that held without the addition no longer holds. But this is just to say that implication may be non-monotonic. But if it's simply accounted for in terms of non-monotonicity then we ought to understand such patterns of implication substructurally (rather than in terms of e.g. defeaters).

Next, I tried to articulate what this approach consisted in. That is, I tried to show what it means to understand the substructurality of implication as inhering in the content itself. In order to do this, I examined various substructural logics, which I believe failed to take the substructurality of implication seriously. I charged these accounts with an implicit commitment to what I called the *assumption of structurality.* This is the assumption that a substructural pattern of implication must explained via the postulation of some further layer of content which stands in fully structural consequence relations. For example, if $A, A, \Gamma \vdash \Delta$ but $A, \Gamma \not\vdash \Delta$ (a violation of contraction), then we must postulate some further contents A_1 and A_2 that A may express. If we then understand this sentece in terms of those contents, there is no violation of contraction should the following hold:

$$A_1, A_2, \Gamma \vdash \Delta$$

$$A_i, \Gamma \not\vdash \Delta \qquad (i \in \{1, 2\})$$

While the sentence A does stand in a non-contractive implication relation, this was explained via the postulation of a level of content that does contract. If all substructural behavior is so understood, then an account is (on my charge) committed to the assumption of structurality and fails to take the substructurality of implication seriously by *pushing it all the way down into the content* of a sentence.

An additional insight that was gained by examining substructural logics was an aprecation of the relationship between substructural approaches to paradox. In particular, I argued that the multi-facetedness of content (which was central for understanding violations of contraction) should play a primary explanatory role in explaining substructural solutions to paradox. Acknowledging this was shown to be central to non-contractive, non-transitive, and non-reflexive solutions to paradox. The disagreement between non-contractive, on the one hand, and non-contractive and non-reflexive solutions, on the other, consisted in whether such facets were inferentially productive. Another insight gained concerned the relationship between non-transitive and non-reflexive solutions to paradox. It has been long acknowledged that transitivity and reflexivity seem to be two sides of the same coin.[1] I clarify the denomination of this coin.

[1](Girard, 1995)

Finally, I showed that this account of substructural content is not only of interest to inferentialists and logicians, but can also do real work in other areas of philosophy. In particular, I showed that appreciating the substructurality of content opens up logical space in the semantic minimalism and radical contextualism debate as well as in the moral particularism and moral generalism debate. In both debates, I showed that a common presupposition underlying the debate seemed to commit the participants to a view of content wherein that content must be involved in fully structural relations of implication. I was also able to show a deep commonality between these debates. Both debates assumed that something like a fully structural principle was needed to appreciate the significance that a consideration has to the content of a sentence, on the one hand, or the normative verdicts concerning action, on the other. I argued that we need not understand the significance of a consideration in this way. If we have a view of content on which the content itself may stand in relations of implication that are substructural, then we can also gain a view of the significance of a consideration under which that significance is understood substructurally and in terms of the contribution it makes to good implication, on the one hand, or normative verdicts, on the other.

I believe that the material developed in this dissertation can be used to develop a rich research project. I hope to leverage the idea of substructural content to not only make substantive debates to other areas of philosophy (in addition to a number of areas in the philosophy of language—vagueness, ambiguity, etc.—I believe that it might be developed into contributions in other areas of meta-ethics (in particular, parts of the reasons literature) as well as epistemology (developing e.g. a substructural account of justification)). In addition to this, there are a number of technical areas that could be further developed. I am interested in not only looking at additional structural rules not explored in the dissertation—permutation is the most obvious, for which there is already lots of material in the philosophy of language concerning orderring effects, but also associativity and other structural concerning sequent-structure—but also expanding the formal semantics to predicates and terms and to whether an account can be given of quantification. The account of substructural content that I've pursued has not only made a large number of contributions within the dissertation, but I hope that it will make a large number in the future as well.

Appendix Unabridged Technical Results

Early in this dissertation I introduce an "implicational phase space semantics". I prove some results concerning this early on and refer back to those results (as well as introduce new results) throughout this dissertation. The results introduced in the dissertation often feature abbreviated proofs (meant to give the shape of a proof without taking up too many pages). The point of this appendix is to collect all of the machinery and proofs together into one location and for the proofs to be presented in an unabridged fashion. Because all of the material appears elsewhere in the dissertation and because I assume that only the most enthusiastic reader will consult this appendix, I am rather sparse on philosophical details here.

Throughout this paper I have introduced in a piecemeal fashion some technical machinery. I have also in a piecemeal fashion and without proper justification claimed certain results concerning this apparatus. In this section I systematically explain how the machinery works independently of its philosophical application above. The following appendices should therefore also serve as a kind of reference for locating features of that system.

A.1 Proof Theory

In this section I rehearse the sequent calculus.

Definition A.1.1 (Base Language and Consequence Relation)**.** A base language $\mathcal{L}_0$ consists of a set of atomic sentence letters: $p_1, \ldots, p_n$.

A base consequence relation relates multi-sets of atoms to multi-sets of atoms:[1]

$$\vdash_0 \subseteq \mathcal{P}(\mathcal{L}_0)^2.$$

We may express such a relation as follows:

$$p_1, \ldots, p_n \vdash_0 q_1, \ldots, q_m,$$

I will typically omit set bracket and union symbols where lack of ambiguity allows. As a convention I will use capital Greek latters $(\Gamma, \Delta, \ldots)$ to represent multi-sets of sentences, capital Latin letters to represent sentences and lowercase Latin letters for atomic sentences (i.e. sentences contained in $\mathcal{L}_0$). I next introduce several important conditions base consequence relations may satisfy.

Definition A.1.2 (Conditions on Base Consequence Relations)**.** The following conditions are of special interest:

Reflexivity A base consequence relation is said to be **reflexive** iff

$$\forall p \in \mathcal{L}_0 (p \vdash_0 p).$$

Monotonicity A base consequence relation is said to be **monotonic** iff[2]

$$\Gamma \vdash_0 \Theta \Rightarrow \forall \Delta, \Lambda \subseteq \mathcal{L}_0 (\Gamma, \Delta \vdash_0 \Theta, \Lambda).$$

[1] Note that the object language consists of "multi-sets", i.e. sets distinguished by multiplicity. While $\{a\}$ and $\{a, a\}$ are the same set, they are *not* the same *multi*-set. For the most part, I don't rely on anything too controversial surrounding multi-sets, but this allows for less ambiguity surrounding how the sequent calculus works.

Note in particular, that I sometimes use notation that is technically ambiguous. For example, I below write that $\vdash_0 \subseteq \mathcal{P}(\mathcal{L}_0)^2$ even though $\vdash_0$ is a relation between multi-sets and $\mathcal{L}_0$ is a set. We should understand $\mathcal{P}(\mathcal{L}_0)$ in this context as the **set** of every possible finite **subset** of $\mathcal{L}_0$ with every possible multiplicity of its members, so that $\mathcal{P}(\mathcal{L}_0)$ is actually a **set** of **multi-sets**, and $\mathcal{P}(\mathcal{L}_0)^2$ is a **set** of ordered pairs of **multi-sets**. In general context shall settle any possible ambiguity here.

[2] Note that I use '$\Rightarrow$', '$\Leftarrow$', and '$\Leftrightarrow$' to signal "meta-theoretic entailment/conditionals". That is $A \Rightarrow B$ means that if A is the case, then B must be the case (via ordinary mathematical/logical reasoning in the meta-theory). I include this note to avoid any potential confusion regarding notation.

Contractivity A base consequence relation is said to be **contractive** iff

$$p, p, \Gamma \vdash_0 \Theta \Rightarrow p, \Gamma \vdash_0 \Theta,$$

and

$$\Gamma \vdash_0 \Theta, p, p \Rightarrow \Gamma \vdash_0 \Theta, p.$$

Flatness A base consequence relation is said to be **flat** iff there is a subset of $\vdash_0$ that is reflexive, monotonic, and contractive. That is, iff

$$\forall p \in \mathcal{L}_0 \forall \Delta, \Lambda \subseteq \mathcal{L}_0(\Delta, p \vdash_0 p, \Lambda).$$

Minimal, Flat $\vdash_0 \subseteq \mathcal{P}(\mathcal{L}_0)^2$ is said to be a **minimal, flat** base consequence relation iff $\vdash_0$ is flat and no proper subset of $\vdash_0$ is also flat.

A.1.1 The Extension

The base language is extended in the standard fashion to include $\{\neg, \vee, \&, \rightarrow\}$:

- If $p \in \mathcal{L}_0$ then $p \in \mathcal{L}$. I call all such 'p' atoms.
- If $A, B \in \mathcal{L}$ then $\ulcorner A \rightarrow B \urcorner, \ulcorner A \& B \urcorner, \ulcorner A \vee B \urcorner \in \mathcal{L}$.
- If $A \in \mathcal{L}$, then $\ulcorner \neg A \urcorner \in \mathcal{L}$.

The base consequence relation ($\vdash_0$) is extended to $\vdash \subseteq \mathcal{P}(\mathcal{L})^2$ according to the following sequent calculus. The axioms, which are the only sequents which may serve as roots of our proof trees are generated as follows:

Axiom If $\Gamma \vdash_0 \Theta$, then $\Gamma \vdash \Theta$.

The rules follow below. I use the Ketonen rules for sentential logic, also sometimes called "assorted" rules (since they mix additive and multiplicative connectives) in the sequent calculus.[3] $\Gamma \vdash \Theta$ is included in $\vdash$ iff there is a deduction from the axioms to $\Gamma \vdash \Theta$ in accordance with these rules:

[3]See Ketonen (1944). See Dicher (2016) for this usage of "assorted" (and also confer Humberstone (2007)). In the Linear Logic, these are sometimes also called "asynchronous" (as opposed to "synchronous") rules. See Andreoli (1992); Liang and Miller (2009). The rules and sequent calculus is also extremely similar to e.g. G3cp, see Negri et al. (2008).

$$\frac{\Gamma \vdash \Theta, A \qquad B, \Gamma \vdash \Theta}{A \rightarrow B, \Gamma \vdash \Theta}\ \mathrm{L}\!\rightarrow \qquad \frac{A, \Gamma \vdash \Theta, B}{\Gamma \vdash A \rightarrow B, \Theta}\ \mathrm{R}\!\rightarrow$$

$$\frac{\Gamma, A, B \vdash \Theta}{\Gamma, A\&B \vdash \Theta}\ \mathrm{L\&} \qquad \frac{\Gamma \vdash \Theta, A \qquad \Gamma \vdash \Theta, B}{\Gamma \vdash \Theta, A\&B}\ \mathrm{R\&}$$

$$\frac{A, \Gamma \vdash \Theta \qquad B, \Gamma \vdash \Theta}{A \vee B, \Gamma \vdash \Theta}\ \mathrm{L}\vee \qquad \frac{\Gamma \vdash \Theta, A, B}{\Gamma \vdash \Theta, A \vee B}\ \mathrm{R}\vee$$

$$\frac{\Gamma \vdash \Theta, A}{\neg A, \Gamma \vdash \Theta}\ \mathrm{L}\neg \qquad \frac{A, \Gamma \vdash \Theta}{\Gamma \vdash \Theta, \neg A}\ \mathrm{R}\neg$$

A.1.2 Some Central Results

First a well known result that will be extremely useful.

Theorem A.1.3. *All of the rules of the sequent calculus are reversible. That is, if* $\Gamma \vdash \Theta$ *would be the result of a rule application to* $\Gamma^* \vdash \Theta^*$ *(and possibly* $\Gamma^{**} \vdash \Theta^{**}$*):*

$$\frac{\Gamma^* \vdash \Theta^* \qquad [\Gamma^{**} \vdash \Theta^{**}]}{\Gamma \vdash \Theta}\ *$$

It follows that

$$\Gamma \vdash \Theta \Leftrightarrow \Gamma^* \vdash \Theta^* (\textit{and possibly } \Gamma^{**} \vdash \Theta^{**}).$$

Proof. By induction on proof height. If $\Gamma \vdash \Theta$ is the result of a proof of height 1 (i.e. only one rule application) then the result is immediate. Therefore suppose the result holds for proof trees of height less than n and suppose $\Gamma \vdash \Theta$ is the result of a proof tree of height n. We must show that the result holds regardless of which rule '*' is. I show the result holds when '*' is L$\rightarrow$ (from which the cases where '*' is L$\vee$ or R& may be proven in an analogous fashion) and when '*' is R$\rightarrow$ (from which the cases where '*' is L&, R$\vee$, L$\neg$, or R$\neg$ may be proven in an analogous fashion.

There are two sub-cases: the last step is either the result of a rule with one-top sequent (R$\rightarrow$, L&, R$\vee$, L$\neg$, R$\neg$) or the result of a rule with two-top sequents (L$\leftarrow$, L$\vee$, R&). We must show that our result holds for all of our rules.

Case 1.1: Suppose '*' is L$\rightarrow$ and the last step is the application of a rule with exactly one top-sequent. Then the last step is of the form:

$$\frac{A \to B, \Gamma' \vdash \Theta'}{A \to B, \Gamma \vdash \Theta}$$

By our inductive hypothesis we have $B, \Gamma' \vdash \Theta'$ and $\Gamma' \vdash \Theta', A$. Via two applications of the same rule we have: $B, \Gamma \vdash \Theta$ and $\Gamma \vdash \Theta, A$.

Case 1.2: Suppose '*' is L$\to$ and the last step is the application of a rule with two top-sequents. If the last step is an application of L$\to$ to our principal formula, then we are done, so suppose this is not the case, then the last step is of the form:

$$\frac{A \to B, \Gamma' \vdash \Theta' \qquad A \to B, \Gamma'' \vdash \Theta''}{A \to B, \Gamma \vdash \Theta}$$

By our inductive hypothesis, we therefore have $B, \Gamma' \vdash \Theta'$ and $\Gamma' \vdash \Theta', A$ as well as $B, \Gamma'' \vdash \Theta''$ and $\Gamma'' \vdash \Theta'', A$. By two applications of the same rule we have $B, \Gamma \vdash \Theta$ and $\Gamma \vdash \Theta, A$.

Case 2.1: Suppose '*' is R$\to$ and the last step is the application of a rule with exactly one top-sequent. If the last step is an application of R$\to$ to our principal formula, then we are sone, so suppose this is not the case, then the last step is of the form:

$$\frac{\Gamma' \vdash \Theta', A \to B}{\Gamma \vdash \Theta, A \to B}$$

By our inductive hypothesis we have $A, \Gamma' \vdash \Theta', B$ and by an application of the same rule we have $A, \Gamma \vdash \Theta, B$.

Case 2.2: Suppose '*' is R$\to$ and the last step is the application of a rule with two top-sequents. Then the last step is of the form:

$$\frac{\Gamma' \vdash \Theta', A \to B \qquad \Gamma'' \vdash \Theta'', A \to B}{\Gamma \vdash \Theta, A \to B}$$

By our inductive hypothesis we have $A, \Gamma' \vdash \Theta', B$ and $A, \Gamma'' \vdash \Theta'', B$. Via an application of the same rule we obtain $A, \Gamma \vdash \Theta, B$.

Cases 3.1,2–8.1,2 Follow in a similar fashion (i.e. when '*' is any of the other rules).

■

The following theorem is also of importance.

Theorem A.1.4. *For every sentence A there exists $A_0^P \subseteq \mathcal{P}(\mathcal{L}_0)^2$ such that*

$$(\forall \langle \Delta, \Lambda \rangle \in A_0^P) A, \Gamma \vdash \Theta \Leftrightarrow \Delta, \Gamma \vdash \Theta, \Lambda.$$

Likewise for every sentence A there exists $A_0^C \subseteq \mathcal{P}(\mathcal{L}_0)^2$ such that

$$(\forall \langle \Delta, \Lambda \rangle \in A_0^P)\Gamma \vdash \Theta, A \Leftrightarrow \Delta, \Gamma \vdash \Theta, \Lambda.$$

These sets (of multi-sets) can be thought of as atomic representatives of A (hence the subscripted 0*; the P and C represent whether A appears in the premise/conclusion of an implication).*

Proof. I prove both halves of the theorem at once by induction on the complexity of A. That is, I show simultaneously the existence of A_0^P and A_0^C.

Suppose A is atomic, then $A_0^P = \{\langle A, \emptyset \rangle\}$ and $A_0^C = \{\langle \emptyset, A \rangle\}$.

Now suppose our result holds for sentences of complexity strictly less than n and suppose A is of complexity n. A may be a conditional, conjunction, disjunction, or negated sentence.

Suppose A is a conditional, then it is of the form $D \to E$. By hypothesis we have $D_0^P, D_0^C, E_0^P, E_0^C$. Clearly $A_0^P = D_0^C \cap E_0^P$ since we have (from our previous theorem) that $D \to E, \Gamma \vdash \Theta$ iff $E, \Gamma \vdash \Theta$ and $\Gamma \vdash \Theta, D$.

Likewise, $A_0^C = \{\{\Delta \cup \Delta', \Lambda \cup \Lambda'\} | \langle \Delta, \Lambda \rangle \in D_0^P$ and $\langle \Delta', \Lambda' \rangle \in E_0^C\}$, since $\Gamma, D \vdash E, \Theta$ iff $\Gamma \vdash D \to E, \Theta$ (so we need the pairwise combination of those sets that comprise D_0^P and E_0^C.

If A is a conjunction or disjunction then we have our result in an analogous manner (consider that $A \& B$ is equivalent to $\neg(A \to B)$ and $A \vee B$ is equivalent to $\neg A \to B$ given our sequent rules).

If A is a negated sentence of the form $\neg B$, then we have by hypothesis B_0^P and B_0^C and so $A_0^P = B_0^C$ and $A_0^C = B_0^P$. ■

We should think of the sets generated in this manner as representing A atomically. That is, they tell us exactly which combination of sets as atoms are needed to build A up in arbitrary contexts (i.e. arbitrary sequents) as either a premise or a conclusion.

From this result, the next result is nearly immediate. First, a definition:

Definition A.1.5 (DNF,CNF)**.** Call a sentence a **literal** if it is either an atom or has the form $\neg A$ where A is an atom.

We say that a sentence A is in **disjunctive normal form (DNF)** if it is a disjunction of conjunctions of literals. I write $DNF(A)$ as shorthand for "the disjunctive normal form of A", where this is understood to mean a sentence that is equivalent to A and in DNF.

We say that a sentence A is in **conjunction normal form (CNF)** if it is a conjunction of disjunctions of literals. I write $CNF(A)$ as shorthand for "the conjunctive normal form of A", where this is understood to mean a sentence that is equivalent to A and in CNF.

Note that in the previous definition and following proposition the expression "equivalent" (sentence) and thus $DNF(A)$ and $CNF(A)$ are ambiguous/not well-defined. The following proposition should therefore be read as establishing the existence of a sentence in DNF that holds exactly when A does, and *mutatis mutandis* for the proposition that follows thereafter.

Proposition A.1.6.

$$A, \Gamma \vdash \Theta \Leftrightarrow DNF(A), \Gamma \vdash \Theta$$

where $DNF(A)$ is the DNF of A.

Proof. Because all rules are reversible, we can simply unfold A in line with Theorem A.1.4 until we have $A_0^P \subseteq \mathcal{P}(\mathcal{L}_0)^2$. Next enumerate A_0^P as $\langle \Delta_1, \Lambda_1 \rangle, \ldots, \langle \Delta_n, \Lambda_n \rangle$. We have

$$A, \Gamma \vdash \Theta \Leftrightarrow \Delta_i, \Gamma \vdash \Theta, \Lambda_i,$$

for $1 \leq i \leq n$. Next let $\neg\Lambda$ be shorthand for $\{\neg\lambda | \lambda \in \Lambda\}$ and $\&\Delta$ be shorthand for $\delta_1 \& \cdots \& \delta_j$ where $\{\delta_1, \ldots, \delta_j\} = \Delta$. We should understand $\bigvee\Delta$ in an analogous fashion. Then we may derive:

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\Delta_1, \Gamma \vdash \Theta, \Lambda_1}{\vdots}\,{\scriptstyle \mathrm{L}\neg}}{\Delta_1, \neg\Lambda_1, \Gamma \vdash \Theta}\,{\scriptstyle \mathrm{L}\neg}}{\vdots}\,{\scriptstyle \mathrm{L}\&}}{\&(\Delta_1 \cup \neg\Lambda_1), \Gamma \vdash \Theta}\,{\scriptstyle \mathrm{L}\&} \qquad \dfrac{\dfrac{\dfrac{\dfrac{\Delta_2, \Gamma \vdash \Theta, \Lambda_2}{\vdots}\,{\scriptstyle \mathrm{L}\neg}}{\Delta_2, \neg\Lambda_2, \Gamma \vdash \Theta}\,{\scriptstyle \mathrm{L}\neg}}{\vdots}\,{\scriptstyle \mathrm{L}\&}}{\&(\Delta_2 \cup \neg\Lambda_2), \Gamma \vdash \Theta}\,{\scriptstyle \mathrm{L}\&}}{\&(\Delta_1 \cup \neg\Lambda_1) \vee \&(\Delta_2 \cup \neg\Lambda_2), \Gamma \vdash \Theta}\,{\scriptstyle \mathrm{L}\vee} \quad \vdots}{\ddots} \qquad \dfrac{\vdots \quad \dfrac{\dfrac{\dfrac{\dfrac{\Delta_n, \Gamma \vdash \Theta, \Lambda_n}{\vdots}\,{\scriptstyle \mathrm{L}\neg}}{\Delta_n, \neg\Lambda_n, \Gamma \vdash \Theta}\,{\scriptstyle \mathrm{L}\neg}}{\vdots}\,{\scriptstyle \mathrm{L}\&}}{\&(\Delta_n \cup \neg\Lambda_n), \Gamma \vdash \Theta}\,{\scriptstyle \mathrm{L}\&}}{\iddots}}{\&(\Delta_1 \cup \neg\Lambda_1) \vee \cdots \vee \&(\Delta_n \cup \neg\Lambda_n), \Gamma \vdash \Theta}\,{\scriptstyle \mathrm{L}\vee}$$

Clearly

$$\&(\Delta_1 \cup \neg\Lambda_1) \vee \cdots \vee \&(\Delta_n \cup \neg\Lambda_n) = \bigvee\{\&(\Delta_i \cup \neg\Lambda_i) | 1 \leq i \leq n\} = DNF(A).$$

■

We may prove a similar result for the CNF in the succedent.

Proposition A.1.7.

$$\Gamma \vdash \Theta, B \Leftrightarrow \Gamma \vdash \Theta, CNF(B)$$

where $CNF(B)$ is the CNF of B.

Proof. Proof is analogous to proof of previous proposition. We shift all the atomic sets (in i.e. B_0^C) to the right-hand side (using R$\neg$), combine them into disjunctions (using R$\vee$) and finally into conjunctions (using R&) yielding a sentence in CNF. ■

Corollary A.1.8. *For arbitrary Γ, Θ, there is some A and some B such that:*

$$\Gamma \vdash \Theta \Leftrightarrow DNF(A) \vdash CNF(B).$$

Proof. Via repeated applications of L& and R$\vee$ we obtain

$$\&\Gamma \vdash \bigvee\Theta.$$

From the previous two propositions this yields

$$DNF(\&\Gamma) \vdash CNF(\bigvee\Theta).$$

So we are therefore guaranteed the existence of some A (in this case $\&\Gamma$) and some B (in this case $\bigvee\Theta$) such that

$$\Gamma \vdash \Theta \Leftrightarrow DNF(A) \vdash CNF(B).$$

■

Definition A.1.9 (Normal Form)**.** We say that a sequent $\Gamma \vdash \Theta$ is in **Normal Form** if its premise consists of a single sentence in DNF, and its conclusion consists of a single sentence in CNF.

A.1.3 Representation Theorem

Finally I prove my representation theorem. In particular I show constructively that given any theory $\mathcal{T}$ expressible in $\mathcal{L}$ we may specify a base consequence relation $\vdash_0$ such that $\vdash$ realizes $\mathcal{T}$. Further I show that if $\mathcal{T}$ meets certain constraints then we may specify a base consequence relation such that $\vdash$ realizes $\mathcal{T}$ *exactly.*

Definition A.1.10 (Theory)**.** We say that $\mathcal{T}$ is a **theory** of sentential logic if $\mathcal{T} \subseteq \mathcal{L}$. We say that $\vdash$ **realizes** $\mathcal{T}$ iff

$$\forall \tau \in \mathcal{T}(\vdash \tau).$$

We say that $\vdash$ realizes $\mathcal{T}$ **exactly** iff

$$\vdash \tau \Leftrightarrow \tau \in \mathcal{T}.$$

Definition A.1.11 (Inferential Relation)**.** We say that $\mathcal{I}$ is an **inferential relation** if $\mathcal{I} \subseteq \mathcal{P}(\mathcal{L})^2$. We say that $\vdash$ **realizes** $\mathcal{I}$ iff

$$\forall \langle \Gamma, \Theta \rangle \in \mathcal{I}(\Gamma \vdash \Theta).$$

We say that $\vdash$ realizes $\mathcal{I}$ **exactly** iff

$$\Gamma \vdash \Theta \Leftrightarrow \langle \Gamma, \Theta \rangle \in \mathcal{I}.$$

First I show that for any arbitrary theory we may find an inferential relation such that they are always co-realized. Likewise for any arbitrary inferential relation we may find a theory such that they are co-realized. Co-satisfaction here means that the one is realized iff the other is realized. Note that "realization" is always relative to a base-consequence relation.

Proposition A.1.12. Let $\mathcal{T} \subseteq \mathcal{L}$. Then we may find $\mathcal{I} \subseteq \mathcal{P}(\mathcal{L})^2$ such that $\mathcal{I}$ is realized iff $\mathcal{T}$ is realized.

Likewise let $\mathcal{I} \subseteq \mathcal{P}(\mathcal{L})^2$, then we may find $\mathcal{T} \subseteq \mathcal{L}$ such that $\mathcal{T}$ is realized iff $\mathcal{I}$ is realized.

Proof. Proof is constructive. Let

$$\mathcal{I} = \{\langle \emptyset, \tau \rangle | \tau \in \mathcal{T}\}.$$

Result is immediate.

Next suppose we have $\mathcal{I}$. Let

$$\mathcal{T} = \{(\&\Gamma) \to (\bigvee \Theta) | \langle \Gamma, \Theta \rangle \in \mathcal{I}\}.$$

■

We have therefore established that realization of theories and realization of inferential relations coincide. But this is a rather weak notion of realization. After all, if we simply let $\vdash_0 = \mathcal{P}(\mathcal{L}_0)^2$ then any theory or inferential relation will be realized vacuously. Next I introduce the notion of a "proper" theory and base consequence relation. I shall prove that this can be realized exactly.

Definition A.1.13 (Proper Theory)**.** We say that a theory is **proper** if it is closed according to the following three rules and four substitution properties.[4]

R1: &-Composition If $\ulcorner A \urcorner, \ulcorner B \urcorner \in \mathcal{T}$ then $\ulcorner A \& B \urcorner \in \mathcal{T}$.

R2: &-Decomposition If $\ulcorner A \& B \urcorner \in \mathcal{T}$ then $\ulcorner A \urcorner, \ulcorner B \urcorner \in \mathcal{T}$.

R3: Distribution $\ulcorner A \vee (B \& C) \urcorner \in \mathcal{T}$ iff $\ulcorner (A \vee B) \& (A \vee C) \urcorner \in \mathcal{T}$.

S1: Conditional Let σ be an instance of a sub-formula of τ. Then $\tau \in \mathcal{T}$ for $\sigma = \ulcorner A \to B \urcorner$ iff $\tau[\sigma/\sigma'] \in \mathcal{T}$ for $\sigma' = \ulcorner \neg A \vee B \urcorner$.

S2: De Morgan's-1 Let σ be an instance of a sub-formula of τ. Then $\tau \in \mathcal{T}$ for $\sigma = \ulcorner \neg (A \vee B) \urcorner$ iff $\tau[\sigma/\sigma'] \in \mathcal{T}$ for $\sigma' = \ulcorner \neg A \& \neg B \urcorner$.

S3: De Morgan's-2 Let σ be an instance of a sub-formula of τ. Then $\tau \in \mathcal{T}$ for $\sigma = \ulcorner \neg (A \& B) \urcorner$ iff $\tau[\sigma/\sigma'] \in \mathcal{T}$ for $\sigma' = \ulcorner \neg A \vee \neg B \urcorner$.

S4: Involution Let σ be an instance of a sub-formula of τ. Then $\tau \in \mathcal{T}$ for $\sigma = \ulcorner \neg\neg B \urcorner$ iff $\tau[\sigma/\sigma'] \in \mathcal{T}$ for $\sigma' = \ulcorner B \urcorner$.

[4] I don't introduce explicitly—though a stricter treatment would require me to—that, for example, theories are required to obey commutivity for '&' and '∨'.

Note that normally when "theories" in the above sense (in sentential or first-order logic) are introduced they are usually introduced not only as a set of sentences, but a set of sentences which is closed in some sense. Usually they are thought to be closed under, for example, classical consequence. All of the above properties are entailed by classical consequence, but note that the criteria for proper theory I have given is strictly *weaker* than classical logic. That is, not all proper theories are theories of classical logic. A symptom of this discrepancy is that proper theories are closed under neither *disjunctive syllogism* nor *modus ponens*.[5]

Example A.1.14. Suppose $\ulcorner A \rightarrow B \urcorner, \ulcorner A \urcorner \in \mathcal{T}$. There is no guarantee that $\ulcorner B \urcorner \in \mathcal{T}$. The only rule that could yield B is R2, but there is no rule or substitution property that yields a sentence with B as a conjunct (the only possibilities are R1 (which would require an independent way to get B) or S2).

Next I show that "proper" theories have an interesting property: they contain the conjunctive normal form of all of their sentences. First a lemma.

Lemma A.1.15. *Suppose $\mathcal{T}$ is proper. Then if $\tau \in \mathcal{T}$ an equivalent sentence σ is also in $\mathcal{T}$ such that σ is either a conjunction or a disjunction.*[6]

Proof. Suppose $\tau \in \mathcal{T}$. If τ is a disjunction or conjunction we are done. It remains that τ may be a conditional or negated sentence. I proceed by induction on complexity. Suppose our result holds for sentences of complexity strictly less than n. Now suppose τ is of complexity n. If τ is a conditional then it is of the form $A \rightarrow B$. By S1 $\ulcorner A \rightarrow B \urcorner \in \mathcal{T}$ iff $\ulcorner \neg A \vee B \urcorner \in \mathcal{T}$. If τ is a negated sentence is of the form $\neg A$. By hypothis if $\ulcorner A \urcorner \in \mathcal{T}$ then there is an equivalent disjunction or conjunction in $\mathcal{T}$. If it is a disjunction of the form $B \vee C$ then by S3 $\ulcorner \neg B \& \neg C \urcorner \in \mathcal{T}$. If it is a conjunction of the form $B \& C$ then by S2 $\ulcorner \neg B \vee \neg C \urcorner \in \mathcal{T}$. ∎

Lemma A.1.16. *Suppose $\mathcal{T}$ is proper and $\tau \in \mathcal{T}$. Then we may find an equivalent sentence $\sigma \in \mathcal{T}$ such that σ is either a disjunction of literals or has '&' as its main connective.*

Proof. I proceed by induction on the complexity of τ. By our previous lemma we need only consider cases in which τ is a conjunction or disjunction. If τ is atomic we are done, so

[5] Note that given S1 and S4, disjunctive syllogism and modus ponens are equivalent.

[6] Note that here "equivalent" means that proper theories contain both or none of these sentences. That is, σ and τ are equivalent iff for an arbitrary theory $\mathcal{T}$: $\sigma \in \mathcal{T}$ iff $\tau \in \mathcal{T}$.

assume the result holds for τ of logical complexity less than n and now let τ be of complexity n. If τ is a conjunction we are done again, so τ let τ be a disjunction of the form $\tau = \ulcorner A \vee B \urcorner$. By hypothesis there is a ρ and ρ' equivalent to A and B, respectively which are either the disjunction of literals or have '&' as their main connective. Suppose ρ and ρ' are both disjunctions of literals, then $\sigma = \ulcorner \rho \vee \rho' \urcorner \in \mathcal{T}$ (clearly the disjunction of disjunctions of literals is also a disjunction of literals).

Next, suppose instead without loss of generality that ρ has '&' as its main connective, i.e. $\rho = \ulcorner C \& D \urcorner$. Then $\ulcorner A \vee (C \& D) \urcorner \in \mathcal{T}$. By R3 we are guaranteed $\ulcorner (A \vee C) \& (A \vee D) \urcorner \in \mathcal{T}$. So let $\sigma = \ulcorner (A \vee C) \& (A \vee D) \urcorner$. ■

Theorem A.1.17. *Let $\mathcal{T}$ be a proper theory. Then*

$$\tau \in \mathcal{T} \Leftrightarrow CNF(\tau) \in \mathcal{T}.$$

Proof. Let $\mathcal{T}$ a proper theory. We show $\tau \in \mathcal{T} \Leftrightarrow CNF(\tau) \in \mathcal{T}$ by induction on logical complexity of τ.

In the base case τ is an atom and so $\tau = CNF(\tau)$. Therefore suppose that our result holds for τ of logical complexity strictly less than n and now suppose that τ is of complexity n. By our previous lemma we need only consider disjunctions of literals and sentences that have '&' as their main connective.

Suppose τ is a conjunction of the form $A \& B$. By R2 we have $\ulcorner A \urcorner, \ulcorner B \urcorner \in \mathcal{T}$ and thus by our inductive hypothesis $CNF(A), CNF(B) \in \mathcal{T}$. By R1 we have $\ulcorner CNF(A) \& CNF(B) \urcorner \in \mathcal{T}$. Clearly $\ulcorner CNF(A) \& CNF(B) \urcorner$ is in conjunctive normal form.

It remains then that τ is a disjunction of literals. In this case τ is already in CNF. ■

What this previous theorem in fact provides is a better characterization of proper theories. We might better understand proper theories as theories which obey R1, R2, and which contains a sentence if and only if it contains the CNF of that sentence.

Corollary A.1.18. *$\mathcal{T}$ is proper iff $\mathcal{T}$ satisfies R1, R2, and the following condition:*

$$\tau \in \mathcal{T} \Leftrightarrow CNF(\tau) \in \mathcal{T}.$$

Definition A.1.19 (Representative)**.** Suppose that $\vdash$ realizes inferential relation $\mathcal{I}$ exactly. Then we may call $\vdash_0$ (i.e. the base consequence relation of $\vdash$) the **representative** of inferential relation $\mathcal{I}$.

Suppose that $\vdash$ realizes $\mathcal{T}$ exactly. Let $\mathcal{I}$ be an equivalent inferential relation and thus $\vdash_0$ its representative. We call the following the **representative** of $\mathcal{T}$:

$$\{\neg p_1 \vee \cdots \vee \neg p_n \vee q_1 \vee \cdots \vee q_m \in \mathcal{L} | p_1, \ldots, p_n \vdash_0 q_1, \ldots, q_m\}.$$

I next demonstrate exactly how these "representatives" are constructed. For each $\langle \Gamma, \Theta \rangle \in \mathcal{I}$ we may determine exactly which atomic implications are required to realize $\Gamma \vdash \Theta$. Via Corollary A.1.8 we have $\Gamma \vdash \Theta$ iff $DNF(A) \vdash CNF(A)$, i.e.

$$\&(\Gamma_1 \cup \neg\Theta_1) \vee \cdots \vee \&(\Gamma_n \cup \neg\Theta_n) \vdash \bigvee(\neg\Delta_1 \cup \Lambda_1) \& \cdots \& \bigvee(\neg\Delta_m \cup \Lambda_m).$$

It follows that (for $1 \leq i \leq n$ and $1 \leq j \leq m$) that:

$$\Gamma_i, \Delta_j \vdash_0 \Theta_i, \Lambda_j.$$

If we take the union of all such sets of atomic implications (for each member of $\mathcal{I}$) then we have our representative.

Since proper $\mathcal{T}$ are characterized by the CNF of each of its members, we may construct its representative in the same manner.

Theorem A.1.20 (Representation Theorem 1)**.** *Let $\mathcal{T}$ (or $\mathcal{I}$) be arbitrary. Then we may specify $\vdash_0$ such that $\vdash$ realizes $\mathcal{T}$ (or $\mathcal{I}$).*

I actually show something stronger. Namely, I show how to expand $\mathcal{T}$ (and likewise $\mathcal{I}$) into minimal, proper theories. I also, thereby, show how to find the smallest $\vdash_0$ that realizes a theory (when it fails to realize that theory exactly). However, in order to minimize the amount of work done here, I prove the second representation theorem first, from which the first follows as a special case.

Theorem A.1.21 (Representation Theorem 2)**.** *Let $\mathcal{T}$ be proper (or $\mathcal{I}$ be equivalent to a proper $\mathcal{T}$). Then we may specify $\vdash_0$ such that $\vdash$ realizes $\mathcal{T}$ (or $\mathcal{I}$)* exactly*.*

Put otherwise: $\mathcal{T}$ may be realized exactly iff *$\mathcal{T}$ is proper.*

Proof. ($\Leftarrow$) Suppose $\mathcal{T}$ is proper. We must find $\vdash_0$ such that $\tau \in \mathcal{T}$ iff $\vdash \tau$.

Let us construct $\vdash_0$ according to the above procedure outline for generating a representative for $\mathcal{T}$.

Let $\vdash_0$ be specified according to the above procedure. For each member of σ of $\mathcal{T}$ with

$$CNF(\sigma) = \bigvee(\neg\Delta_1 \cup \Lambda_1)\& \cdots \& \bigvee(\neg\Delta_m \cup \Lambda_m),$$

we stipulate that $\vdash_0$ contain

$$\Delta_i \vdash \Lambda_j,$$

for $1 \leq i \leq m$ and $1 \leq j \leq m$. Next, given that $\tau \in \mathcal{T}$ iff $CNF(\tau) \in \mathcal{T}$ (via Theorem A.1.17), clearly $\vdash CNF(\tau)$ since we have stipulated that $\vdash_0$ contains the materials for constructing $\vdash CNF(\tau)$—and I note that this connection is biconditional. Further, I note that $\vdash CNF(\tau)$ iff $\vdash \tau$ (via Proposition A.1.7). It therefore follows that

$$\tau \in \mathcal{T} \Leftrightarrow \vdash \tau.$$

That is $\vdash$ realizes $\mathcal{T}$ *exactly.*

($\Rightarrow$) Let $\mathcal{T}$ be a theory and suppose $\vdash$ realizes $\mathcal{T}$ *exactly.* We wish to show that $\mathcal{T}$ is proper. Via corollary A.1.18 it is sufficient to show that $\mathcal{T}$ obeys R1, R2 and that

$$\tau \in \mathcal{T} \Leftrightarrow CNF(\tau) \in \mathcal{T}.$$

The last stipulation is immediate since we have shown (via Proposition A.1.7) that

$$\vdash \tau \Leftrightarrow \vdash CNF(\tau).$$

Next, I note that:

$$\vdash A\&B \Leftrightarrow \vdash A \text{ and } \vdash B,$$

It follows that

$$\ulcorner A\&B \urcorner \in \mathcal{T} \Leftrightarrow A, B \in \mathcal{T}.$$

That is, $\mathcal{T}$ satisfies R1 and R2. ■

Now, I return to a proof of the first representation theorem. I first, however, introduce the following definition, which streamlines proof of the result and is of independent interest.

Definition A.1.22 (Minimal, Proper Theory)**.** Let $\mathcal{T}$ be an arbitrary theory. Let $CNF(\mathcal{T})$ specify the set that contains all and only the CNF sentences of every sentence in $\mathcal{T}$:

$$CNF(\mathcal{T}) = \{CNF(\tau)|\tau \in \mathcal{T}\}.$$

We call $\mathcal{S}$ the **Minimal, Proper Theory** of $\mathcal{T}$ iff $\mathcal{S}$ is proper and

$$CNF(A) \in \mathcal{S} \Leftrightarrow CNF(A) \in CNF(\mathcal{T}).$$

We can use a similar procedure to define the **Minimal, Proper Consequence Relation** of a consequence relation $\mathcal{I}$.

Now, here is the proof of the first Representation Theorem (A.1.20).

Proof. Let $\mathcal{T}$ be an arbitrary theory. Let $\tau \in \mathcal{T}$ be arbitrary. We wish to find $\vdash_0$ such that $\vdash \tau$.

Next let $\mathcal{S}$ be the *minimal, proper theory* of $\mathcal{T}$. It follows that $CNF(\tau) \in \mathcal{S}$. Now, from the proof of A.1.21 we are guaranteed the existence of $\vdash_0$ and $\vdash$ that realize $\mathcal{S}$ exactly. That is: $\vdash CNF(\tau)$. Further (via Theorem A.1.17) we have that $\tau \in \mathcal{S}$.[7] Thus $\vdash \tau$. ∎

Corollary A.1.23. *Let $\mathcal{T}$ be an arbitrary theory, let $\mathcal{S}$ be its minimal, proper theory, and let $\vdash_0$ be such that $\vdash$ realizes $\mathcal{S}$* exactly*. It follows that $\vdash_0$ (and $\vdash$) is minimal with respect to consequence relations that satisfy $\mathcal{T}$.*

[7] Alternatively, via Proposition A.1.7, we have the next step directly.

A.1.3.1 Applications

I delay a more thorough discussion of the application of the representation theorem for later, but I would like to foreshadow now its use. What the representation theorems tell us is that a specification of atomic behavior is not only sufficient (but necessary) for specifying the behavior of logically complex implications. This means that given a theory, we may find the behavior of any of the sentences that appears within it and specify that sentence's role in terms of $\vdash_0$. If we think theories capable of implicitly defining the sentences that occur within them, then this is a powerful result for connecting the theory to a notion of meaning (i.e. role in implication).

Further, what the notion of representative (Definition A.1.19) gives us, is a way of specifying for proper theories, the material element of a consequence relation, that is, the relation between atomic sentences that must obtain for a theory (or consequence relation) to obtain.

A.1.4 Theories

I next explain ways in which we may limit what counts as a proper $\vdash_0$. At the start I introduced several constraints we may wish to place on $\vdash_0$, e.g. contractivity, reflexivity, monotonicity, flatness, etc. I explain some interesting features that result from these constraints as well explain how to implement so-far undiscussed constraints. In particular I aim to the prove the following results:[8]

1. $\vdash_0$ is monotonic iff $\vdash$ is monotonic.
2. $\vdash_0$ is contractive iff $\vdash$ is contractive.
3. $\vdash_0$ is flat iff $\vdash$ is flat.
4. Minimal, flat $\vdash_0$ is equivalent to classical logic (i.e. $\vdash_{LK}$).
5. Existence of a condition on $\vdash_0$ for which $\vdash$ is transitive (i.e. admits cut-elimination).
6. Theories: application of representation theorem

I prove these results in order. First a lemma.

[8](1),(3), and (4) are already well-established, though I go through the proofs here for the sake of completeness. Most of the machinery to establish (2) and (5) exists independently of my apparatus, though I am unaware of a discussion of exactly these results.

Lemma A.1.24. *We may weaken a sequent with arbitrary atoms iff we may weaken it with arbitrary sentences (of potential logical complexity):*

$$\forall \Delta_0, \Lambda_0 \subseteq \mathcal{L}_0(\Delta_0, \Gamma \vdash \Theta, \Lambda_0) \Leftrightarrow \forall \Delta, \Lambda \subseteq \mathcal{L}(\Delta, \Gamma \vdash \Theta, \Lambda).$$

Proof. The ($\Leftarrow$)-direction is immediate. The ($\Rightarrow$)-direction is also proven fairly quickly. Let Δ and Λ be arbitrary. We have that:[9]

$$\Delta, \Gamma \vdash \Theta, \Lambda \Leftrightarrow \bigwith \Delta, \Gamma \vdash \Theta, \bigvee \Lambda.$$

We therefore have

$$DNF(\bigwith \Delta), \Gamma \vdash \Theta, CNF(\bigvee \Lambda).$$

But this is just equivalent to

$$\bigwith(\Gamma_1 \cup \neg\Theta_1) \vee \cdots \vee \bigwith(\Gamma_n \cup \neg\Theta_n), \Gamma \vdash \Theta, \bigvee(\neg\Delta_1 \cup \Lambda_1) \& \cdots \& \bigvee(\neg\Delta_m \cup \Lambda_m).$$

It follows that (for $1 \leq i \leq n$ and $1 \leq j \leq m$) that:

$$\Gamma_i, \Delta_j, \Gamma \vdash_0 \Theta, \Theta_i, \Lambda_j.$$

Since $\Gamma_i, \Delta_j, \Theta_i, \Lambda_j$ are atomic (for $1 \leq i \leq n$ and $1 \leq j \leq m$), we have our result. ∎

Theorem A.1.25. $\vdash_0$ *is monotonic iff* $\vdash$ *is monotonic.*

Proof. The ($\Leftarrow$) direction is immediate. The ($\Rightarrow$) direction follows via induction on proof height.

Let Δ, Λ be arbitrary. We must show $\Delta, \Gamma \vdash \Theta, \Lambda$ (given that $\vdash_0$ is atomic). Now, in the base case $\Gamma \vdash \Theta$ is atomic. Via Lemma A.1.24, we have our result.

Now suppose our result holds for proof trees of height strictly less than n and that $\Gamma \vdash \Theta$ is obtained via a proof tree of height n. It may come via any of our rules. I divide into two cases where $\Gamma \vdash \Theta$ comes either via a rule with one-top sequent or via a rule with two-top sequents.

Case 1: Suppose $\Gamma \vdash \Theta$ comes via a rule with one-top sequent:

[9] An independent proof of result—that doesn't take advantage of my previously introduced machinery—can be constructed via induction on logical complexity.

$$\frac{\Gamma^* \vdash \Theta^*}{\Gamma \vdash \Theta} *$$

By hypothesis we have $\Delta, \Gamma^* \vdash \Theta^*, \Lambda$. Via an application of the same rule (*) we have our result: $\Delta, \Gamma \vdash \Theta, \Lambda$.

Case 2: Suppose $\Gamma \vdash \Theta$ comes via a rule with two-top sequents:

$$\frac{\Gamma^* \vdash \Theta^* \qquad \Gamma^{**} \vdash \Theta^{**}}{\Gamma \vdash \Theta} **$$

By hypothesis we have $\Delta, \Gamma^* \vdash \Theta^*, \Lambda$ and $\Delta, \Gamma^{**} \vdash \Theta^{**}, \Lambda$. Via an application of the same rule (**) we have our result: $\Delta, \Gamma \vdash \Theta, \Lambda$.

■

Next, I show that a similar result holds for contractivity.

Lemma A.1.26. *Suppose $\vdash$ obeys contraction of atoms. I.e. $p, p, \Gamma \vdash \Theta$ only if $p, \Gamma \vdash \Theta$ and $\Gamma \vdash \Theta, p, p$ only if $\Gamma \vdash \Theta, p$. It follows that $\vdash$ obeys contraction of arbitrary sentences, i.e. $A, A, \Gamma \vdash \Theta$ only if $A, \Gamma \vdash \Theta$ and $\Gamma \vdash \Theta, A, A$ only if $\Gamma \vdash \Theta, A$.*

Proof. I proceed via induction on the complexity of A. In the base case A is atomic and we have our result immediately. Therefore suppose our result holds when A is of logical complexity strictly less than n and now suppose A is of complexity n. A may be a conditional, conjunction, disjunction, or negated sentence.

If A is a conditional then it is of the form $B \to C$. Then if we have $B \to C, B \to C, \Gamma \vdash \Theta$, we must have:

$$\Gamma \vdash \Theta, B, B \tag{1}$$

$$C, \Gamma \vdash \Theta, B$$

$$C, C, \Gamma \vdash \Theta \tag{2}$$

Applying our inductive hypothesis to (1) and (2) yields

$$\Gamma \vdash \Theta, B$$

$$C, \Gamma \vdash \Theta$$

Via L$\to$ we have $B \to C, \Gamma \vdash \Theta$.

Likewise if we have $\Gamma \vdash \Theta, B \to C, B \to C$ then this must come from $B, B, \Gamma \vdash \Theta, C, C$. Via our inductive hypothesis we have $B, \Gamma \vdash \Theta, C$ and via R$\to$ we have $\Gamma \vdash \Theta, B \to C$.

The cases where A is either a conjunction or disjunction are handled analogously.

If A is a negation, i.e. of the form $\neg B$, then we may reason as follows. If we have $\neg B, \neg B, \Gamma \vdash \Theta$ then this must come via $\Gamma \vdash \Theta, B, B$. By hypothesis $\Gamma \vdash \Theta, B$ and via L$\neg$ we have $\neg B, \Gamma \vdash \Theta$. If we have $\Gamma \vdash \Theta, \neg B, \neg B$ then $\Gamma \vdash \Theta, \neg B$ follows in a similar fashion. ■

Theorem A.1.27. $\vdash_0$ *is contractive iff* $\vdash$ *is contractive.*

Proof. The ($\Leftarrow$) direction is immediate so I show the ($\Rightarrow$) direction. It is sufficient to show that contraction of atoms is preserved (i.e. that $\vdash$ allows contraction of atomic sentences). Lemma A.1.26 secures the result after that.

I proceed therefore via induction on proof height. In the base case $\Gamma \vdash \Theta$ is atomic, so we are done. Suppose, therefore, that $\Gamma \vdash \Theta$ admits contraction of atoms for proof trees of height strictly less than n and that $p, p, \Gamma \vdash \Theta$ and $\Gamma \vdash \Theta, p, p$ each come via proof trees of height n. They may come via any of our rules. Because p is atomic we can be certain that it is not the principal formula in our rule application. I treat two cases (where they come via a rule with one top sequent and where they come via a rule with two top sequents).

Case 1: Suppose $p, p, \Gamma \vdash \Theta$ (and $\Gamma \vdash \Theta, p, p$) come via a rule with one-top sequent:

$$\frac{p, p, \Gamma^* \vdash \Theta^* \text{ (or } \Gamma^* \vdash \Theta^*, p, p)}{\Gamma \vdash \Theta \text{ (or } \Gamma \vdash \Theta, p, p)} *$$

By hypothesis we have $p, \Gamma^* \vdash \Theta^*$ (and $\Gamma^* \vdash \Theta^*, p$). Via an application of the same rule (*) we have our result: $p, \Gamma \vdash \Theta$ (and $\Gamma \vdash \Theta, p$).

Case 2: Suppose $p, p, \Gamma \vdash \Theta$ (and $\Gamma \vdash \Theta, p, p$) come via a rule with two-top sequents:

$$\frac{p, p, \Gamma^* \vdash \Theta^* \text{ (or } \Gamma^* \vdash \Theta^*, p, p) \qquad p, p, \Gamma^{**} \vdash \Theta^{**} \text{ (or } \Gamma^{**} \vdash \Theta^{**}, p, p)}{p, p, \Gamma \vdash \Theta \text{ (or } \Gamma \vdash \Theta, p, p)} **$$

By hypothesis we have $p, \Gamma^* \vdash \Theta^*$ (and $\Gamma^* \vdash \Theta^*, p$) as well as $p, \Gamma^{**} \vdash \Theta^{**}$ (and $\Gamma^{**} \vdash \Theta^{**}, p$). Via an application of the same rule (**) we have our result: $p, \Gamma \vdash \Theta$ (and $\Gamma \vdash \Theta, p$).

■

Next, I show that $\vdash_0$ is flat iff $\vdash$ is flat, i.e.

$$\forall \Gamma_0, \Theta_0, p(\Gamma_0, p \vdash_0 p, \Theta_0) \Leftrightarrow \forall \Gamma, \Theta, A(\Gamma, A \vdash A, \Theta).$$

I have already shown that monotonicity is preserved so it remains to show that reflexivity in preserved in such contexts.

Lemma A.1.28. *If $\vdash_0$ is flat, then $\vdash$ is reflexive, i.e. $A \vdash A$ for all A.*

Proof. I proceed via induction on the complexity of A. If A is atomic we are done, so suppose the result holds when A is of complexity strictly less than n and now suppose A is of complexity n. It may be a conditional, conjunction, disjunction, or negated sentence.

Suppose A is a conditional of the form $B \to C$. By hypothesis we have $B \vdash C, B$ and $B, C \vdash C$. We derive:

$$\dfrac{\dfrac{C, B \vdash C \qquad B \vdash C, B}{B \to C, B \vdash C}\,\text{L}\!\to}{B \to C \vdash B \to C}\,\text{R}\!\to$$

$\vee$ and & are handled analogously. Therefore suppose that A is an negated sentence of the form $\neg B$. By hypothesis we have $B \vdash B$ and so

$$\dfrac{\dfrac{B \vdash B}{\neg B, B \vdash}\,\text{L}\neg}{\neg B \vdash \neg B}\,\text{R}\neg$$

∎

The result is immediate.

Theorem A.1.29. $\vdash_0$ *is flat iff* $\vdash$ *is flat.*

Proof. Since flat $\vdash_0$ gives us reflexive $\vdash$ and such sequents will be monotonic we are done. ∎

We have so far constructed nearly all of the machinery needed to demonstrate that minimal, flat $\vdash_0$ is equivalent to classical logic. In fact the result should be obvious given the following facts:

- Minimal, flat $\vdash_0$ obeys monotonicity, reflexivity, and *contraction* (i.e. these structural rules are admissible).
- Given monotonicity and contraction, the rules of my sequent calculus are equivalent to e.g. Gentzen's LK rules ($\vdash_{LK}$).

- In $\vdash_{LK}$ the sole axiom is idempotence (i.e. reflexivity). Since we are guaranteed $A \vdash A$ for all A it should be clear that we can reconstruct all such axioms.

Since all of these are well-attested (and in general the equivalence to LK is well-known) I will not spend much time proving this. Instead I'll briefly sketch for the reader unfamiliar with these results how the rules of of my sequent calculus and LK are equivalent given minimal, flat $\vdash_0$. First, I list here all of Gentzen's rules (I replace some of his symbols to keep these two systems distinct).

$$\frac{\Delta \vdash_{LK} \Lambda, A \qquad B, \Gamma \vdash_{LK} \Theta}{A \supset B, \Gamma \vdash_{LK} \Theta}\ \text{L}\supset \qquad \frac{A, \Gamma \vdash_{LK} \Theta, B}{\Gamma \vdash_{LK} A \supset B, \Theta}\ \text{R}\supset$$

$$\frac{\Gamma, A_i \vdash_{LK} \Theta}{\Gamma, A_1 \wedge A_2 \vdash_{LK} \Theta}\ \text{L}\wedge\ (i = 1 \text{ or } 2) \qquad \frac{\Delta \vdash_{LK} \Lambda, A \qquad \Gamma \vdash_{LK} \Theta, B}{\Gamma, \Delta \vdash_{LK} \Theta, \Lambda, A \wedge B}\ \text{R}\wedge$$

$$\frac{A, \Delta \vdash_{LK} \Lambda \qquad B, \Gamma \vdash_{LK} \Theta}{A \vee_{LK} B, \Gamma, \Delta \vdash_{LK} \Theta, \Lambda}\ \text{L}\vee_{LK} \qquad \frac{\Gamma \vdash_{LK} \Theta, A_i}{\Gamma \vdash_{LK} \Theta, A_1 \vee_{LK} A_2}\ \text{R}\vee_{LK}\ (i = 1 \text{ or } 2)$$

$$\frac{\Gamma \vdash_{LK} \Theta, A}{{\sim}A, \Gamma \vdash_{LK} \Theta}\ \text{L}{\sim} \qquad \frac{A, \Gamma \vdash_{LK} \Theta}{\Gamma \vdash_{LK} \Theta, {\sim}A}\ \text{R}{\sim}$$

Gentzen also includes the structural rules (I don't list permutation):

$$\frac{\Gamma \vdash_{LK} \Theta}{\Delta, \Gamma \vdash_{LK} \Theta, \Lambda}\ \text{MO} \qquad \frac{A, A, \Gamma \vdash_{LK} \Theta}{A, \Gamma \vdash_{LK} \Theta}\ \text{L-Contr.} \qquad \frac{\Gamma \vdash_{LK} \Theta, A, A}{\Gamma \vdash_{LK} \Theta, A}\ \text{R-Contr.}$$

And the axiom:

$$\frac{}{A \vdash_{LK} A}\ \text{Idem.}$$

It should be obvious how all proof trees of LK can be replicated. We can construct any axioms of LK given a minimal, flat $\vdash_0$. Further, any application of a structural rule (MO or Contraction) are guaranteed from the results in this section. It remains that we can find equivalent rule applications. I show why this is the case for conjunction (the other cases are either analogous or immediate).

I show any valid application of L$\wedge$ may be reproduced in my sequent calculus and vice-versa. Suppose we have:

$$\frac{\Gamma, A \vdash_{LK} \Theta}{\Gamma, A \wedge B \vdash_{LK} \Theta}\ \text{L}\wedge$$

In my sequent calculus, if we have $\Gamma, A \vdash \Theta$ then we also have $\Gamma, A, B \vdash \Theta$. Via L& we have $\Gamma, A\&B \vdash \Theta$.

Conversely, if we have:

$$\frac{\Gamma, A, B \vdash \Theta}{\Gamma, A\&B, \vdash \Theta}\text{ L\&}$$

We may find an equivalent derivation in LK:

$$\cfrac{\cfrac{\cfrac{\Gamma, A, B \vdash_{LK} \Theta}{\Gamma, A \wedge B, B \vdash_{LK} \Theta}\text{ L}\wedge}{\Gamma, A \wedge B, A \wedge B \vdash_{LK} \Theta}\text{ L}\wedge}{\Gamma, A \wedge B \vdash_{LK} \Theta}\text{ L-Contr.}$$

We may show something similar for conjunction in the succedent. Suppose we have:

$$\frac{\Delta \vdash_{LK} \Lambda, A \qquad \Gamma \vdash_{LK} \Theta, B}{\Gamma, \Delta \vdash_{LK} \Theta, \Lambda, A \wedge B}\text{ R}\wedge$$

Then if we have $\Delta \vdash \Lambda, A$ and $\Gamma \vdash \Theta, \Lambda, B$ we also have $\Delta, \Gamma \vdash \Lambda, A$ and $\Delta, \Gamma \vdash \Theta, \Lambda, B$ (since $\vdash_0$ is monotonic). Thus via R& we have: $\Delta, \Gamma \vdash \Theta, \Lambda, A\&B$.

Conversely, if we have:

$$\frac{\Gamma \vdash \Theta, A \qquad \Gamma \vdash \Theta, B}{\Gamma \vdash \Theta, A\&B}\text{ R\&}$$

We may find an equivalent derivation in LK:

$$\cfrac{\cfrac{\cfrac{\cfrac{\cfrac{\Gamma \vdash_{LK} \Theta, A \qquad \Gamma \vdash_{LK} \Theta, B}{\Gamma, \Gamma \vdash_{LK} \Theta, \Theta, A \wedge B}\text{ R}\wedge}{\vdots}\text{ L-Contr.}}{\Gamma \vdash_{LK} \Theta, \Theta, A \wedge B}\text{ L-Contr.}}{\vdots}\text{ R-Contr.}}{\Gamma \vdash_{LK} \Theta, A \wedge B}\text{ R-Contr.}$$

From these considerations the following result is immediate:

Theorem A.1.30. *$\vdash_0$ is flat iff $\vdash$ is* **supra-classical** *in the sense that any implication of classical logic is included in $\vdash$, i.e. $\vdash_{LK} \subseteq \vdash$.*

This only demonstrates a rather limited version of supra-classicality (that the logic include all classically valid implications). Often, when we wish to characterize a consequence relation as supra-classical, we have something stronger in mind: that it obeys classical principles of reasoning.[10] That is, that $\vdash$ be (at least) monotonic and transitive. It is possible,

[10]Parenthetically: it is not so clear that one can cleanly separate these two things. But I do not address this here.

for example, for $\vdash$ to be supraclassical (and so $\vdash_0$ to be flat) even though $A \vdash B$ is not monotonic (if it occurs outside the fragment of $\vdash$ generated from the flat portion of $\vdash_0$). Likewise if we also have $B \vdash C$ (supposing this also occurs outside this fragment) we might have no guarantee that $A \vdash C$. To many, to character this as a supra-classical would be misleading (some might even refuse to call such a $\vdash$ a consequence relation at all). It will therefore be useful to explain how we can limit $\vdash_0$ to produce a *transitive* $\vdash$, i.e. a $\vdash$ that obeys cut.

Cut is typically formulated as:

$$\frac{\Delta \vdash \Lambda, A \qquad A, \Gamma \vdash \Theta}{\Delta, \Gamma \vdash \Theta, \Lambda}\,\text{Cut}$$

It is easy to see, however, that together with a flat $\vdash_0$ imposing such a condition will result in monotonicity. It is easy to see why. Let Δ, Λ be arbitrary. Suppose $A \vdash B$. Clearly $A, \Delta \vdash \Lambda, A$. Via Cut we have:

$$\frac{A, \Delta \vdash \Lambda, A \qquad A \vdash B}{\Delta, A \vdash B, \Lambda}\,\text{Cut}$$

Therefore, in order to keep these ideas notions separate it will be useful to examine instead a more restricted version of Cut:

$$\frac{\Gamma \vdash \Theta, A \qquad A, \Gamma \vdash \Theta}{\Gamma \vdash \Theta}\,\text{Shared-Cut}$$

Given MO and Contraction these two rules are obviously equivalent (again), so if one insists on monotonic $\vdash_0$ this choice makes no large difference.

It remains then to show what $\vdash_0$ must look like for this structural feature to hold. From previous proofs it should be apparent that we will be done if we find a condition such that: if $A, \Gamma_0 \vdash \Theta_0$ and $\Gamma_0 \vdash \Theta_0, A$ for atomic Γ_0, Θ_0, then $\Gamma_0 \vdash \Theta_0$.

The semantics I introduce in the next section make it significantly easier to formulate this condition, nevertheless I give a sense of what it looks like now. What we want is that anytime the atomic sequents used to construct $A, \Gamma \vdash \Theta$ and $\Gamma \vdash \Theta, A$ appear, that $\Gamma \vdash_0 \Theta$. But how do we specify such behavior? We can think of it as follows. Let A be some sentence, then anytime $DNF(A), \Gamma \vdash \Theta$ and $\Gamma \vdash \Theta, CNF(A)$, we have $\Gamma \vdash_0 \Theta$.

Recall that $DNF(A), \Gamma \vdash \Theta$ is generated by sequents with the following shape:

$$\Delta_1, \Gamma \vdash_0 \Theta, \Lambda_1$$
$$\vdots$$
$$\Delta_n, \Gamma \vdash_0 \Theta, \Lambda_n.$$

We can think of these as ordered pairs which, when combined with $\Gamma \vdash_0 \Theta$, generate a valid sequent. Thus, let us represent $DNF(A)$ on the left with the following set:

$$\{\langle \Delta_1, \Lambda_1 \rangle, \ldots, \langle \Delta_n, \Lambda_n \rangle\}.$$

Then we may represent $CNF(A)$ on the right as:

$$\{\langle \Pi, \Xi \rangle | \delta_i \in \Delta_i, \lambda_i \in \Lambda_i \text{ and } \textit{exactly} \text{ one of either } \delta_i \in \Xi \text{ or } \lambda_i \in \Pi \text{ (for } 1 \leq i \leq n)\}.$$

I.e. take one element from exactly one of the sets in each $\langle \Delta_i, \Lambda_i \rangle$ to construct each $\langle \Pi, \Xi \rangle$. If we enumerate this set as $\langle \Pi_1, \Xi_1 \rangle, \ldots, \langle \Pi_m, \Xi_m \rangle$, then $\Gamma \vdash \Theta, CNF(A)$ will come via:

$$\Pi_1, \Gamma \vdash_0 \Theta, \Xi_1$$
$$\vdots$$
$$\Pi_m, \Gamma \vdash_0 \Theta, \Xi_m.$$

Definition A.1.31 (Transitivity)**.** Let $A \subseteq \mathcal{P}(\mathcal{L}_0)$ be

$$A = \{\langle \Delta_1, \Lambda_1 \rangle, \ldots, \langle \Delta_n, \Lambda_n \rangle\},$$

and let $S \subseteq \mathcal{P}(\mathcal{L}_0)$ be

$$S = \{\langle \Pi, \Xi \rangle | \langle \Delta_i, \Lambda_i \rangle \in A, \delta_i \in \Delta_i, \lambda_i \in \Lambda_i \text{ and } \textit{exactly} \text{ one of either } \delta_i \in \Xi \text{ or } \lambda_i \in \Pi \text{ (for } 1 \leq i \leq n)\}.$$

Let us enumerate S as $\langle \Pi_1, \Xi_1 \rangle, \ldots, \langle \Pi_m, \Xi_m \rangle$. Let Γ, Θ be arbitrary. If whenever

$$\Delta_1, \Gamma \vdash_0 \Theta, \Lambda_1$$
$$\vdots$$
$$\Delta_n, \Gamma \vdash_0 \Theta, \Lambda_n.$$

and

$$\Pi_1, \Gamma \vdash_0 \Theta, \Xi_1$$
$$\vdots$$
$$\Pi_m, \Gamma \vdash_0 \Theta, \Xi_m.$$

then

$$\Gamma \vdash_0 \Theta,$$

then we say that $\vdash_0$ is **transitive** or satisfies **transitivity**.

Lemma A.1.32. *Suppose $\vdash_0$ satisfies transitivity. Then whenever $\Gamma_0, \Theta_0 \subseteq \mathcal{L}_0$ are such that $A, \Gamma_0 \vdash \Theta_0$ and $\Gamma_0 \vdash \Theta_0, A$ we have $\Gamma_0 \vdash \Theta_0$.*

Proof. Immediate from definition. Let A be in DNF which has the form:

$$\&(\Delta_1 \cup \neg\Lambda_1) \vee \cdots \vee \&(\Delta_n \cup \neg\Lambda_n), \Gamma_0 \vdash \Theta_0.$$

Clearly this comes via

$$\Delta_1, \Gamma \vdash_0 \Theta, \Lambda_1$$
$$\vdots$$
$$\Delta_n, \Gamma \vdash_0 \Theta, \Lambda_n.$$

in $\vdash_0$. Next, by supposing we have $\Gamma_0 \vdash \Theta_0, A$, i.e.

$$\Gamma_0 \vdash \Theta_0, \&(\Delta_1 \cup \neg\Lambda_1) \vee \cdots \vee \&(\Delta_n \cup \neg\Lambda_n).$$

I show how we can use this to construct S as in Definition A.1.31.

Since our rules are reversible, we therefore have:

$$\Gamma_0 \vdash \Theta_0, \&(\Delta_1 \cup \neg\Lambda_1), \ldots, \&(\Delta_n \cup \neg\Lambda_n).$$

Let us enumerate each Δ_i as $\delta_{i,1}, \ldots, \delta_{i,k_i}$ and each Λ_i as $\lambda_{i,1}, \ldots, \lambda_{i,k_i}$. If we focus on the final disjunction, then decompositions yields:

$$\begin{aligned}
\Gamma_0 &\vdash \Theta_0, \&(\Delta_1 \cup \neg\Lambda_1), \ldots, \&(\Delta_{n-1} \cup \neg\Lambda_{n-1}), \delta_{n,1} \\
&\vdots \\
\Gamma_0 &\vdash \Theta_0, \&(\Delta_1 \cup \neg\Lambda_1), \ldots, \&(\Delta_{n-1} \cup \neg\Lambda_{n-1}), \delta_{n,k_n} \\
\lambda_{n,1}, \Gamma_0 &\vdash \Theta_0, \&(\Delta_1 \cup \neg\Lambda_1), \ldots, \&(\Delta_{n-1} \cup \neg\Lambda_{n-1}) \\
&\vdots \\
\lambda_{n,k_n}, \Gamma_0 &\vdash \Theta_0, \&(\Delta_1 \cup \neg\Lambda_1), \ldots, \&(\Delta_{n-1} \cup \neg\Lambda_{n-1}).
\end{aligned}$$

Repeating this procedure for each of the above generated sequents, it is clear that we can construct a set

$$S = \{\langle \Pi, \Xi \rangle | \langle \Delta_i, \Lambda_i \rangle \in A, \delta_i \in \Delta_i, \lambda_i \in \Lambda_i \text{ and } \textit{exactly} \text{ one of either } \delta_i \in \Xi \text{ or } \lambda_i \in \Pi \text{ (for } 1 \leq i \leq n)\}.$$

We are therefore done. ■

Theorem A.1.33. $\vdash_0$ *is transitive iff Shared-Cut is admissible:*

$$\dfrac{\Gamma \vdash \Theta, A \qquad A, \Gamma \vdash \Theta}{\Gamma \vdash \Theta}\ \textit{Shared-Cut}$$

Proof. (⇒) Suppose $\vdash_0$ is transitive. We show that cut is admissible by induction on proof height. We suppose that proofs assume a normal form under which the principal formula that is to be cut is constructed first and other rule applications are applied afterwards. In the base case $p, \Gamma \vdash \Theta$ and $\Gamma \vdash \Theta, p$ are atomic so we are done. Therefore suppose that the result holds for proof trees of height strictly less than n. Now suppose that $A, \Gamma \vdash \Theta$ and $\Gamma \vdash \Theta, A$ come from proof trees of at most height n. Regardless of which rule they issue from, we may use our inductive hypothesis to "cut" A and then via an application of the same rule obtain $\Gamma \vdash \Theta$.

(⇐) Immediate from previous lemma. ■

What we have therefore are the tools for restricting $\vdash_0$ as we see fit. We may either restrict it by imposing certain structural conditions on $\vdash_0$ (monotonicity, contraction, flatness, transitivity, etc.) from which we know certain structural features of $\vdash$ to follow. Or we may restrict $\vdash_0$ so that certain theories or inferential relations are realized in $\vdash$. Together these tools will give us powerful means for exploring how, to what degree, and in what sense new sentential variables may be defined in a logic.

A.2 Semantics

In this section I introduce the inferential role semantics and prove that the proof theory (i.e. sequent calculus) of the previous section is *sound* and *complete* with respect to this semantics.

A.2.1 Inferential Spaces

Definition A.2.1 (Inferential Space)**.** Given a language $\mathcal{L}$ (of potential logical complexity), an **inferential space** is understood to be the set of ordered pairs of multi-sets of $\mathcal{L}$:[11]

$$\mathbf{P} = \mathcal{P}(\mathcal{L})^2.$$

We call each "point" (of the form $\langle X, Y\rangle$, where $X, Y \subseteq \mathcal{L}$) an **implication**.

Each inferential space $\mathbf{P}$ comes with a privileged subset of implications: the **good implications**.

$$\mathbb{I} \subseteq \mathbf{P}.$$

NB: I have so far introduced *no constraints* on what $\mathbb{I}$ must look like (nor on what $\mathbf{P}$) must look like. I have left it entirely open how each of these spaces is constituted.[12]

[11] See n.1 for more information on my choice to use multi-sets.

[12] Compare my treatment of $\vdash_0$ in the Appendix A.

Definition A.2.2 (fuission)**.** There is a single operation on **P** known as **fuission**, '$\sqcup$'. If $A = \langle \Gamma, \Theta \rangle$ and $B = \langle \Delta, \Lambda \rangle$, then[13]

$$A \sqcup B =_{df.} \langle \Gamma \cup \Delta, \Theta \cup \Lambda \rangle.$$

I also generalize '$\sqcup$' as an operation over subsets of **P**. If $X, Y \subseteq \mathbf{P}$, then:

$$X \sqcup Y = \{x \cup y | x \in X, y \in Y\}.$$

Definition A.2.3 (vee)**.** I next define a function which will eventually help to specify the inferential role of sentences. '$\curlyvee$' (pronounced "vee") is a function from subsets of **P** to subsets of **P**. '$\curlyvee$' tells us the role an ordered pair makes to good implication. Thus if $X = \langle \Gamma, \Delta \rangle$, then:

$$X^{\curlyvee} = \{\langle \Gamma, \Delta \rangle\}^{\curlyvee} =_{df.} \{\langle \Delta, \Lambda \rangle | \langle \Gamma, \Delta \rangle \sqcup \langle \Delta, \Lambda \rangle \in \mathbb{I}\}.$$

When X is a singleton, we may drop the set brackets where lack of ambiguity allows. In the case where $X \subseteq \mathbf{P}$ and is not a singleton, then '$\curlyvee$' is defined as:

$$X^{\curlyvee} =_{df.} \{\langle \Delta, \Lambda \rangle | \forall \langle \Gamma, \Theta \rangle \in X(\langle \Gamma, \Theta \rangle \sqcup \langle \Delta, \Lambda \rangle \in \mathbb{I})\}.$$

Definition A.2.4 (Closure)**.** '$\curlyvee$' allows us to define a closure operation. A set of implications $X \subseteq \mathbf{P}$ is said to be **closed** iff

$$X^{\curlyvee\curlyvee} = X.$$

I now prove that '$\curlyvee$' is in fact a *closure* operation in the way described. Traditionally *closure operations* $Cl(\cdot)$ are said to satisfy three properties:

Extensive: A closure contains that which it is the closure of: $X \subseteq Cl(X)$.
Idempotent: The closure of a closure is the same as the closure: $Cl(Cl(X)) = Cl(X)$.
Monotone: If $X \subseteq Y$, then $Cl(X) \subseteq Cl(Y)$.

Proposition A.2.5. '$\curlyvee\curlyvee$' is a closure operation.

[13]Note that '$\cup$' is here understood to be multi-set union, i.e. in this context $\{a\} \cup \{a\} = \{a, a\}$.

Proof. Let $X \subseteq \mathbf{P}$ be arbitrary. We must show that '$\curlyvee\curlyvee$' serves as a closure operation on such X.

Clearly '$\curlyvee\curlyvee$' is extensive. Let $\langle \Gamma, \Delta \rangle \in X$ be arbitrary. Now if $\langle x_1, x_2 \rangle \in X^{\curlyvee}$ then $\langle \Gamma, \Delta \rangle \sqcup \langle x_1, x_2 \rangle \in \mathbb{I}$. Since $\langle x_1, x_2 \rangle$ was chosen arbitrarily it follows that $\langle \Gamma, \Delta \rangle \in X^{\curlyvee\curlyvee}$. Thus:

$$X \subseteq X^{\curlyvee\curlyvee}.$$

'$\curlyvee\curlyvee$' is also idempotent. Since we have already shown it to be extensive, we know that $X^{\curlyvee\curlyvee} \subseteq X^{\curlyvee\curlyvee\curlyvee\curlyvee}$. Thus we need only show the converse, i.e. $X^{\curlyvee\curlyvee\curlyvee\curlyvee} \subseteq X^{\curlyvee\curlyvee}$. Let $a \in X^{\curlyvee}$ be arbitrary and $b \in X^{\curlyvee\curlyvee\curlyvee\curlyvee}$ be arbitrary. We must show $a \sqcup b \in \mathbb{I}$. We know that $a \in X^{\curlyvee\curlyvee\curlyvee}$. Thus, $a \sqcup b \in \mathbb{I}$ and so $a \in X^{\curlyvee\curlyvee}$. It follows that:

$$X^{\curlyvee\curlyvee\curlyvee\curlyvee} = X^{\curlyvee\curlyvee}.$$

Finally, '$\curlyvee\curlyvee$' is monotone. Let $X \subseteq Y \subseteq \mathbf{P}$. Now if $a \in Y^{\curlyvee}$ then for arbitrary $b \in Y$ we have that $a \sqcup b \in \mathbb{I}$. It therefore follows that $a \in X^{\curlyvee}$. Now let $c \in X^{\curlyvee\curlyvee}$ be arbitrary. Clearly $a \sqcup c \in \mathbb{I}$ and thus $c \in Y^{\curlyvee\curlyvee}$. It follows that:

$$X^{\curlyvee\curlyvee} \subseteq Y^{\curlyvee\curlyvee}.$$

■

Three corollaries follow immediately.

Corollary A.2.6. $X^{\curlyvee} = X^{\curlyvee\curlyvee\curlyvee}$.

Corollary A.2.7. *If* $X \subseteq Y$ *then* $Y^{\curlyvee} \subseteq X^{\curlyvee}$.

The final corollary is a result of the "extensive" property of closure operations.

Corollary A.2.8. $X \subseteq \mathbb{I}$ *iff* $X^{\curlyvee\curlyvee} \subseteq \mathbb{I}$.

The important or interesting feature of the previous corollary is that $X \in \mathbb{I}$ (i.e. X is a *good* implication) iff $X^{\curlyvee\curlyvee} \subseteq \mathbb{I}$.

Now I introduce the notion of a "proper inferential role".

Definition A.2.9 (Proper Inferential Role)**.** A **proper inferential role (PIR)** is a double $\langle X, Y\rangle$ such that X and Y are each *closed*—in the sense described above—subsets of $\mathbf{P}$ (i.e. $X^{\curlyvee\curlyvee} = X$ and $Y^{\curlyvee\curlyvee} = Y$). We should think of $\langle X, Y\rangle$ as specifying some inferential role and think of X as specifying that premissory role and Y as specifying that conclusory role.

As a convention if $A = \langle X, Y\rangle$ is an inferential role, then I write A^P to refer to X and A^C to refer to Y.

I call these inferential roles "proper" in order to distinguish it from a more informal notion of inferential role which need not be closed in the relevant sense.

Next, I introduce some important semantic notions.

A.2.2 Semantics

Definition A.2.10 (Models)**.** A **model** is a tuple $\langle \mathbf{P}, \mathbb{I}, [\![\cdot]\!]\rangle$ consisting of an language $\mathcal{L}$ and inferential space over that language $\mathbf{P}$, a privileged set of good implications $\mathbb{I}$, and an interpretation function $[\![\cdot]\!]$ (to be defined next) which interprets sentences in the language as inferential roles in the model.

.

Definition A.2.11 (Interpretation Function)**.** An interpretation function $[\![\cdot]\!]$ maps sentences in $\mathcal{L}$ to proper inferential roles in models (i.e. ordered pairs of closed sets of implications). If $A \in \mathcal{L}$ is atomic, then A is interpreted as follows:

$$[\![A]\!] =_{df.} \langle \langle \{A\}, \emptyset\rangle^{\curlyvee}, \langle \emptyset, \{A\}\rangle^{\curlyvee}\rangle.$$

My semantic definitions follow:

$$\begin{aligned}
[\![A\&B]\!] &=_{df.} \langle (([\![A]\!]_P)^{\curlyvee} \sqcup ([\![B]\!]_P)^{\curlyvee})^{\curlyvee}, [\![A]\!]_C \cap [\![B]\!]_C\rangle,\\
[\![A \vee B]\!] &=_{df.} \langle [\![A]\!]_P \cap [\![B]\!]_P, (([\![A]\!]_C)^{\curlyvee} \sqcup ([\![B]\!]_C)^{\curlyvee})^{\curlyvee}\rangle,\\
[\![A \to B]\!] &=_{df.} \langle [\![A]\!]_C \cap [\![B]\!]_P, (([\![A]\!]_P)^{\curlyvee} \sqcup ([\![B]\!]_C)^{\curlyvee})^{\curlyvee}\rangle,\\
[\![\neg A]\!] &=_{df.} \langle [\![A]\!]_C, [\![A]\!]_P\rangle.
\end{aligned}$$

It is a legitimate question whether our semantic definitions successfully pick out *inferential roles* (in the sense made clear in Definition A.2.9). As I've already proven, for arbitrary $X \in \mathbf{P}$, we have that $X^{\curlyvee}$ is closed. This clarifies how many of our semantic definitions yield closed sets. But it is not immediately obvious why $X \cap Y$ will be closed even if both X and Y are. We must therefore provide a short proof of this. First, a useful lemma:

Lemma A.2.12. *Suppose X and Y are closed, then:*

$$X \cap Y = (X^{\curlyvee} \cup Y^{\curlyvee})^{\curlyvee}.$$

Proof. Let $a \in X \cap Y$. Let $b \in X^{\curlyvee} \cup Y^{\curlyvee}$. Either $b \in X^{\curlyvee}$ or $b \in Y^{\curlyvee}$. Without loss of generality, we have that $a \sqcup b \in \mathbf{P}$ (since $a \in X$ and $a \in Y$). It therefore follows that $a \in (X^{\curlyvee} \cup Y^{\curlyvee})^{\curlyvee}$.

Conversely, suppose $a \in (X^{\curlyvee} \cup Y^{\curlyvee})^{\curlyvee}$. Let $b \in (X^{\curlyvee} \cup Y^{\curlyvee})$. Either $b \in X^{\curlyvee}$ or $b \in Y^{\curlyvee}$. In either case $a \sqcup b \in \mathbf{P}$. Thus, $a \in X^{\curlyvee\curlyvee}$ and $a \in Y^{\curlyvee\curlyvee}$. Since X and Y are closed it follows that $a \in X$ and $a \in Y$, i.e. $a \in X \cap Y$. ∎

Proposition A.2.13. If X and Y are closed, then so is $X \cap Y$.

Proof. We must show $(X \cap Y)^{\curlyvee\curlyvee} = X \cap Y$. From our lemma we have that:

$$(X \cap Y)^{\curlyvee\curlyvee} = (X^{\curlyvee} \cup Y^{\curlyvee})^{\curlyvee\curlyvee\curlyvee}.$$

We also know from a previous proof that

$$(X^{\curlyvee} \cup Y^{\curlyvee})^{\curlyvee\curlyvee\curlyvee} = (X^{\curlyvee} \cup Y^{\curlyvee})^{\curlyvee}.$$

Finally, applying the previous lemma again, we have:

$$(X^{\curlyvee} \cup Y^{\curlyvee})^{\curlyvee} = X \cap Y.$$

Thus,

$$(X \cap Y)^{\curlyvee\curlyvee} = X \cap Y.$$

. ∎

We can also prove an interesting feature of our semantic definitions. This requires two lemmas, however, the second of which more or less guarantees the result:

Lemma A.2.14.

$$X^{\curlyvee\curlyvee} \sqcup Y^{\curlyvee\curlyvee} \subseteq (X \sqcup Y)^{\curlyvee\curlyvee}.$$

Proof. Let $a \in X^{\curlyvee\curlyvee}$ and $b \in Y^{\curlyvee\curlyvee}$. Next let $c \in (X \sqcup Y)^{\curlyvee}$. We must show that $a \sqcup b \sqcup c \in \mathbb{I}$.

Next let $x \in X$ and $y \in Y$. It follows that $c \sqcup x \sqcup y \in \mathbb{I}$. It thus follows that $c \sqcup y \in X^{\curlyvee}$. It therefore follows that $a \sqcup c \sqcup y \in \mathbb{I}$. Since $y \in Y$ it follows that $a \sqcup c \in Y^{\curlyvee}$ and thus $a \sqcup b \sqcup c \in \mathbb{I}$. Thus $a \sqcup b \in (X \sqcup Y)^{\curlyvee\curlyvee}$. ■

Before preceding I note that a significant corollary follows from the previous lemma.

Corollary A.2.15. $(X^{\curlyvee\curlyvee} \sqcup Y^{\curlyvee\curlyvee})^{\curlyvee\curlyvee} = (X \sqcup Y)^{\curlyvee\curlyvee}$.

Proof. From the previous lemma we have $X^{\curlyvee\curlyvee} \sqcup Y^{\curlyvee\curlyvee} \subseteq (X \sqcup Y)^{\curlyvee\curlyvee}$, thus

$$(X^{\curlyvee\curlyvee} \sqcup Y^{\curlyvee\curlyvee})^{\curlyvee\curlyvee} \subseteq (X \sqcup Y)^{\curlyvee\curlyvee}.$$

Further, note that since $X \subseteq X^{\curlyvee\curlyvee}$ and $Y \subseteq Y^{\curlyvee\curlyvee}$ clearly $X \sqcup Y \subseteq X^{\curlyvee\curlyvee} \sqcup Y^{\curlyvee\curlyvee}$ and so

$$(X \sqcup Y)^{\curlyvee\curlyvee} \subseteq (X^{\curlyvee\curlyvee} \sqcup Y^{\curlyvee\curlyvee})^{\curlyvee\curlyvee}.$$

■

Lemma A.2.16.

$$(X^{\curlyvee} \sqcup (Y \cap Z)^{\curlyvee})^{\curlyvee} = (X^{\curlyvee} \sqcup Y^{\curlyvee})^{\curlyvee} \cap (X^{\curlyvee} \sqcup Z^{\curlyvee})^{\curlyvee}.$$

Proof. I prove containment in both directions.

($\subseteq$) Suppose $a \in (X^{\curlyvee} \sqcup (Y \cap Z)^{\curlyvee})^{\curlyvee}$. Note that (via Lemma A.2.12) $a \in (X^{\curlyvee} \sqcup (Y^{\curlyvee} \cup Z^{\curlyvee})^{\curlyvee\curlyvee})^{\curlyvee}$. Next let $b \in X^{\curlyvee}$ and let $c \in Y^{\curlyvee}$ or $c \in Z^{\curlyvee}$. Clearly $c \in (Y^{\curlyvee} \cup Z^{\curlyvee})^{\curlyvee\curlyvee}$. It follows that $a \sqcup b \sqcup c \in \mathbb{I}$. Since we left it open whether $c \in^{\curlyvee}$ or $c \in Z^{\curlyvee}$ it follows that $a \in (X^{\curlyvee} \sqcup Y^{\curlyvee})^{\curlyvee}$ and $a \in (X^{\curlyvee} \sqcup Z^{\curlyvee})^{\curlyvee}$. Thus $a \in (X^{\curlyvee} \sqcup Y^{\curlyvee})^{\curlyvee} \cap (X^{\curlyvee} \sqcup Z^{\curlyvee})^{\curlyvee}$.

($\supseteq$) We reason as follows. First note that

$$X^{\curlyvee} \sqcup (Y^{\curlyvee} \cup Z^{\curlyvee}) = (X^{\curlyvee} \sqcup Y^{\curlyvee}) \cup (X^{\curlyvee} \sqcup Z^{\curlyvee}).$$

Next, it follows that

$$\begin{aligned}
X^{\curlyvee\curlyvee\curlyvee} \sqcup (Y^{\curlyvee} \cup Z^{\curlyvee})^{\curlyvee\curlyvee} &= X^{\curlyvee} \sqcup (Y^{\curlyvee} \cup Z^{\curlyvee})^{\curlyvee\curlyvee} \\
&= X^{\curlyvee} \sqcup (Y \cap Z)^{\curlyvee} && \text{(Lemma A.2.12)} \\
&\subseteq ((X^{\curlyvee} \sqcup Y^{\curlyvee}) \cup (X^{\curlyvee} \sqcup Z^{\curlyvee}))^{\curlyvee\curlyvee}. && \text{(Lemma A.2.14)}
\end{aligned}$$

Note next that in general $(X^{\curlyvee} \sqcup Y^{\curlyvee}) \subseteq (X^{\curlyvee} \sqcup Y^{\curlyvee})^{\curlyvee\curlyvee}$ and $(X^{\curlyvee} \sqcup Z^{\curlyvee}) \subseteq (X^{\curlyvee} \sqcup Z^{\curlyvee})^{\curlyvee\curlyvee}$. Thus:

$$\begin{aligned}
((X^{\curlyvee} \sqcup Y^{\curlyvee}) \cup (X^{\curlyvee} \sqcup Z^{\curlyvee}))^{\curlyvee\curlyvee} &\subseteq ((X^{\curlyvee} \sqcup Y^{\curlyvee})^{\curlyvee\curlyvee} \cup (X^{\curlyvee} \sqcup Z^{\curlyvee})^{\curlyvee\curlyvee})^{\curlyvee\curlyvee} \\
&= ((X^{\curlyvee} \sqcup Y^{\curlyvee})^{\curlyvee\curlyvee\curlyvee} \cap (X^{\curlyvee} \sqcup Z^{\curlyvee})^{\curlyvee\curlyvee\curlyvee})^{\curlyvee\curlyvee\curlyvee} && \text{(Lemma A.2.12)} \\
&= ((X^{\curlyvee} \sqcup Y^{\curlyvee})^{\curlyvee} \cap (X^{\curlyvee} \sqcup Z^{\curlyvee})^{\curlyvee})^{\curlyvee}.
\end{aligned}$$

The result so far is thus:

$$X^{\curlyvee} \sqcup (Y \cap Z)^{\curlyvee} \subseteq ((X^{\curlyvee} \sqcup Y^{\curlyvee})^{\curlyvee} \cap (X^{\curlyvee} \sqcup Z^{\curlyvee})^{\curlyvee})^{\curlyvee}.$$

Corollary A.2.7 thus yields

$$((X^{\curlyvee} \sqcup Y^{\curlyvee})^{\curlyvee} \cap (X^{\curlyvee} \sqcup Z^{\curlyvee})^{\curlyvee}) = ((X^{\curlyvee} \sqcup Y^{\curlyvee})^{\curlyvee} \cap (X^{\curlyvee} \sqcup Z^{\curlyvee})^{\curlyvee})^{\curlyvee\curlyvee} \subseteq (X^{\curlyvee} \sqcup (Y \cap Z)^{\curlyvee})^{\curlyvee}.$$

■

Several interesting facts follow:

Corollary A.2.17.

$$\begin{aligned}
[\![A\&(B \vee C)]\!]^P &= [\![(A\&B) \vee (A\&C)]\!]^C \\
[\![A \vee (B\&C)]\!]^C &= [\![(A \vee B)\&(A \vee C)]\!]^C \\
[\![(A\&B) \to C]\!]^P &= [\![(A \to C) \vee (B \to C)]\!]^P \\
[\![(A \vee B) \to C]\!]^C &= [\![(A \to B)\&(A \to C)]\!]^C
\end{aligned}$$

Definition A.2.18 (Semantic Entailment)**.** We say that A semantically entails B relative to a model $\mathcal{M}$ if closure of the fuission of A (as premise) and B (as conclusion) consists of only good implications:

$$A \vDash_{\mathcal{M}} B \quad \text{iff}_{df.} \quad ((\llbracket A\rrbracket^P)^\curlyvee \sqcup (\llbracket B\rrbracket^C)^\curlyvee)^{\curlyvee\curlyvee} \subseteq \mathbb{I}_{\mathcal{M}}.$$

We say that A semantically entails B if $A \vDash_{\mathcal{M}} B$ on all models $\mathcal{M}$.
NB: If A and B are sets of sentences then we read $A \vdash B$ as $\&A \vdash \bigvee B$, i.e. the conjunction of the elements of A and the disjunction of the elements of B.

In order to simplify matter, it is worth noting that the following holds in the semantics.

Proposition A.2.19. Our semantic clause for '&' in the premises and '$\vee$' in the conclusion is associative. In particular:

$$\llbracket (A\&B)\&C\rrbracket^P = \llbracket A\&(B\&C)\rrbracket^P = \llbracket A\&B\&C\rrbracket^P.$$

Put simpler:

$$(((X^\curlyvee \sqcup Y^\curlyvee)^\curlyvee)^\curlyvee \sqcup Z^\curlyvee)^\curlyvee = (X^\curlyvee \sqcup ((Y^\curlyvee \sqcup Z^\curlyvee)^\curlyvee)^\curlyvee)^\curlyvee = (X^\curlyvee \sqcup Y^\curlyvee \sqcup Z^\curlyvee)^\curlyvee.$$

Proof. Since '$\sqcup$' is commutative it will be sufficient to show:

$$(((X^\curlyvee \sqcup Y^\curlyvee)^\curlyvee)^\curlyvee \sqcup Z^\curlyvee)^\curlyvee = (X^\curlyvee \sqcup Y^\curlyvee \sqcup Z^\curlyvee)^\curlyvee.$$

First, we note that $X^\curlyvee \sqcup Y^\curlyvee \subseteq ((X^\curlyvee \sqcup Y^\curlyvee)^\curlyvee)^\curlyvee$. It follows that:

$$X^\curlyvee \sqcup Y^\curlyvee \sqcup Z^\curlyvee \subseteq ((X^\curlyvee \sqcup Y^\curlyvee)^\curlyvee)^\curlyvee \sqcup Z^\curlyvee$$

From Corollary A.2.7, we therefore have that:

$$(((X^\curlyvee \sqcup Y^\curlyvee)^\curlyvee)^\curlyvee \sqcup Z^\curlyvee)^\curlyvee \subseteq (X^\curlyvee \sqcup Y^\curlyvee \sqcup Z^\curlyvee)^\curlyvee.$$

So we need only establish the converse.

Let $a \in (X^\curlyvee \sqcup Y^\curlyvee \sqcup Z^\curlyvee)^\curlyvee$. Next let $x \in X^\curlyvee$, $y \in Y^\curlyvee$ and $z \in Z^\curlyvee$. Clearly $a \sqcup x \sqcup y \sqcup z \in \mathbb{I}$. Since these were all arbitrarily chosen it follows straightaway that $a \sqcup z \in (X^\curlyvee \sqcup Y^\curlyvee)^\curlyvee$. Now suppose $d \in ((X^\curlyvee \sqcup Y^\curlyvee)^\curlyvee)^\curlyvee$. It follows that $a \sqcup z \sqcup d \in \mathbb{I}$. But we may re-arrange this as $d \sqcup z \sqcup a \in \mathbb{I}$, thus

$$a \in (((X^\curlyvee \sqcup Y^\curlyvee)^\curlyvee)^\curlyvee \sqcup Z^\curlyvee)^\curlyvee,$$

and so we are done. ■

The following is immediate:

Corollary A.2.20.

$$\gamma_1, \ldots, \gamma_n \vDash_{\mathcal{M}} \theta_1, \ldots, \theta_m \Leftrightarrow ((\llbracket\gamma_1\rrbracket^P)^{\curlyvee} \sqcup \cdots \sqcup (\llbracket\gamma_n\rrbracket^P)^{\curlyvee} \sqcup (\llbracket\theta_1\rrbracket^C)^{\curlyvee} \sqcup \cdots \sqcup (\llbracket\theta_m\rrbracket^C)^{\curlyvee})^{\curlyvee\curlyvee} \subseteq \mathcal{I}_{\mathcal{M}}.$$

Unfortunately, due to how I've set things up, no entailments will hold across all models. Here is a short proof.

Proof. Suppose $A \vDash B$. For simplicity, suppose both A and B are atomic. This means that

$$(\langle\{A\}, \emptyset\rangle^{\curlyvee\curlyvee\curlyvee\curlyvee} \sqcup \langle\emptyset, \{B\}\rangle^{\curlyvee\curlyvee\curlyvee\curlyvee})^{\curlyvee\curlyvee} \subseteq \mathbb{I}_{\mathcal{M}},$$

for all models $\mathcal{M}$. We may simplify this as:

$$(\langle\{A\}, \emptyset\rangle^{\curlyvee\curlyvee} \sqcup \langle\emptyset, \{B\}\rangle^{\curlyvee\curlyvee})^{\curlyvee\curlyvee} \subseteq \mathbb{I}_{\mathcal{M}},$$

Since $\langle\{A\}, \emptyset\rangle \in \langle\{A\}, \emptyset\rangle^{\curlyvee\curlyvee}$ and $\langle\{A\}, \emptyset\rangle \in \langle\{B\}, \emptyset\rangle^{\curlyvee\curlyvee}$, it follows that $\langle\{A\}, \{B\}\rangle \in (\langle\{A\}, \emptyset\rangle^{\curlyvee\curlyvee} \sqcup \langle\emptyset, \{B\}\rangle^{\curlyvee\curlyvee})^{\curlyvee\curlyvee}$. Thus $\langle\{A\}, \{B\}\rangle \in \mathbb{I}_{\mathcal{M}}$ for all models $\mathcal{M}$. But clearly nothing we have said forces this result. We may easily construct a model in which $\langle\{A\}, \{B\}\rangle \notin \mathbb{I}$. ■

We therefore require a way of restricting the class of models. To do this I (re-)introduce the notion of a base consequence relation. The choice to restrict our attention to atomic consequence relations is (in a certain sense) arbitrary. We could have just as easily restricted out attention to any kind of consequence relation. The choice to do so in this fashion is designed to complement the proof theory as well as the philosophical account given in the first part of this document.

Definition A.2.21 (Base Consequence Relation)**.** A **base consequence relation** is a subset of $\mathbf{P}$ that consists of only atoms. B is a base consequence relation iff $B \subseteq \mathbf{P}$ and $B \cap \mathcal{P}(\mathcal{L}_0)^2 = B$.

We say that a model $\mathcal{M} = \langle \mathbf{P}, \mathbb{I}, \llbracket\cdot\rrbracket\rangle$ is **fit for** a base consequence relation B iff

$$\forall\langle\Delta, \Lambda\rangle \in B(\Delta \vDash_{\mathcal{M}} \Lambda).$$

We say that Γ semantically entails Θ **relative to** B iff $\Gamma \vDash_{\mathcal{M}} \Theta$ for all models $\mathcal{M}$ fit for B.

A.2.3 Soundness and Completeness

I now prove that the semantics is sound and complete with respect to the sequent calculus. To do so we'll need to enrich how we've formulated things there. I therefore introduce the following notation:

$$\Gamma \vdash_B \Theta,$$

To mean that Γ implies Θ **relative to base consequence relation** B. I shall prove that the proof theory is sound and complete in the sense that (for arbitrary Γ, Θ, and B):

$$\Gamma \vdash_B \Theta \Leftrightarrow \Gamma \vDash_B \Theta.$$

Theorem A.2.22 (Soundness)**.** *The sequent calculus is sound:*

$$\Gamma \vdash_B \Theta \Rightarrow \Gamma \vDash_B \Theta.$$

Proof. Let B be arbitrary. I show that if $\Gamma \vdash_B \Theta$ then $\Gamma \vDash_B \Theta$ by induction on proof height.

In the base case $\Gamma \vdash_B \Theta$ is atomic (i.e. $\Gamma, \Theta \subseteq \mathcal{L}_0$). Since we have restricted our attention to models *fit for* B this case it immediate. That is, $\Gamma \vDash_B \Theta$.

Next suppose that our result holds for proof trees of height strictly less than n and now suppose that $\Gamma \vdash_B \Theta$ comes as the last step of a proof tree of height n. It may come from any of the rules of our sequent calculus (L$\to$, R$\to$, L&, R&, L$\vee$, R$\vee$, L$\neg$, R$\neg$).

If the result comes from a rule with a single-top sequent (i.e. R$\to$, L&, R$\vee$, L$\neg$, or R$\neg$), the result is more or less immediate.

Let us enumerate Γ as $\gamma_1, \ldots, \gamma_n$ and Θ as $\theta_1, \ldots, \theta_n$. Our hypothesis can be restated as (by Corollary A.2.20):

$$((\llbracket\gamma_1\rrbracket^P)^\curlyvee \sqcup \cdots \sqcup (\llbracket\gamma_n\rrbracket^P)^\curlyvee \sqcup (\llbracket\theta_1\rrbracket^C)^\curlyvee \sqcup \cdots \sqcup (\llbracket\theta_m\rrbracket^C)^\curlyvee)^{\curlyvee\curlyvee} \subseteq \mathcal{I}_\mathcal{M}.$$

First I note the following equivalences:

$$(((\llbracket\gamma_i\rrbracket^P)^\curlyvee(\llbracket\theta_j\rrbracket^C)^\curlyvee)^\curlyvee = \llbracket\gamma_i \to \theta_j\rrbracket^C$$
$$(((\llbracket\gamma_i\rrbracket^P)^\curlyvee(\llbracket\theta_j\rrbracket^P)^\curlyvee)^\curlyvee = \llbracket\gamma_i \& \gamma_j\rrbracket^P$$
$$(((\llbracket\theta_i\rrbracket^C)^\curlyvee(\llbracket\theta_j\rrbracket^C)^\curlyvee)^\curlyvee = \llbracket\theta_i \vee \theta_j\rrbracket^C$$
$$\llbracket\theta_i\rrbracket^C = \llbracket\neg\theta_i\rrbracket^P$$
$$\llbracket\gamma_i\rrbracket^P = \llbracket\neg\gamma_i\rrbracket^C$$

The following is therefore immediate from our inductive hypothesis:

$$((\llbracket\gamma_2\rrbracket^P)^\curlyvee \sqcup \cdots \sqcup (\llbracket\gamma_n\rrbracket^P)^\curlyvee \sqcup (\llbracket\gamma_1 \to \theta_1\rrbracket^C)^\curlyvee \sqcup \llbracket\theta_2\rrbracket^C)^\curlyvee \sqcup \cdots \sqcup (\llbracket\theta_m\rrbracket^C)^\curlyvee)^{\curlyvee\curlyvee} \subseteq \mathcal{I}_\mathcal{M}$$
$$((\llbracket\gamma_1 \& \gamma_2\rrbracket^P)^\curlyvee \sqcup (\llbracket\gamma_3\rrbracket^P)^\curlyvee \sqcup \cdots \sqcup (\llbracket\gamma_n\rrbracket^P)^\curlyvee \sqcup (\llbracket\theta_1\rrbracket^C)^\curlyvee \sqcup \cdots \sqcup (\llbracket\theta_m\rrbracket^C)^\curlyvee)^{\curlyvee\curlyvee} \subseteq \mathcal{I}_\mathcal{M}$$
$$((\llbracket\gamma_1\rrbracket^P)^\curlyvee \sqcup \cdots \sqcup (\llbracket\gamma_n\rrbracket^P)^\curlyvee \sqcup (\llbracket\theta_1 \vee \theta_2\rrbracket^C)^\curlyvee \sqcup (\llbracket\theta_3\rrbracket^C)^\curlyvee \sqcup \cdots \sqcup (\llbracket\theta_m\rrbracket^C)^\curlyvee)^{\curlyvee\curlyvee} \subseteq \mathcal{I}_\mathcal{M}$$
$$((\llbracket\neg\theta_1\rrbracket^P)^\curlyvee \sqcup (\llbracket\gamma_1\rrbracket^P)^\curlyvee \sqcup \cdots \sqcup (\llbracket\gamma_n\rrbracket^P)^\curlyvee \sqcup (\llbracket\theta_2\rrbracket^C)^\curlyvee \sqcup \cdots \sqcup (\llbracket\theta_m\rrbracket^C)^\curlyvee)^{\curlyvee\curlyvee} \subseteq \mathcal{I}_\mathcal{M}$$
$$((\llbracket\gamma_2\rrbracket^P)^\curlyvee \sqcup \cdots \sqcup (\llbracket\gamma_n\rrbracket^P)^\curlyvee \sqcup (\llbracket\neg\gamma_i\rrbracket^C)^\curlyvee \sqcup (\llbracket\theta_1\rrbracket^C)^\curlyvee \sqcup \cdots \sqcup (\llbracket\theta_m\rrbracket^C)^\curlyvee)^{\curlyvee\curlyvee} \subseteq \mathcal{I}_\mathcal{M}$$

Thus it follows that

$$\Gamma \vDash_B \Theta,$$

whenever $\Gamma \vdash_B \Theta$ is derivable in proof trees of height n whose last step contains only one-top sequent.

We must show the same when $\Gamma \vdash_B \Theta$ comes via one of our rules with two-top sequents (i.e. L$\to$, R&, and L$\vee$). The key here consists in rewriting our semantic entailment clause as a single set-intersection (that we can do this was established in Lemma A.2.16). I nonetheless explain in detail. The last step in our proof trees is thus:

$$B \to A, \gamma_1, \ldots, \gamma_n \vdash_B \theta_1, \ldots, \theta_m$$
$$\gamma_1, \ldots, \gamma_n \vdash_B \theta_1, \ldots, \theta_m, A \& B$$
$$A \vee B, \gamma_1, \ldots, \gamma_n \vdash_B \theta_1, \ldots, \theta_m$$

The top sequents in these cases shall be

$$A, \gamma_1, \ldots, \gamma_n \vdash_B \theta_1, \ldots, \theta_m \quad \text{and} \quad \gamma_1, \ldots, \gamma_n \vdash_B \theta_1, \ldots, \theta_m, B \tag{L$\to$}$$

$$\gamma_1, \ldots, \gamma_n \vdash_B \theta_1, \ldots, \theta_m, A \quad \text{and} \quad \gamma_1, \ldots, \gamma_n \vdash_B \theta_1, \ldots, \theta_m, B \tag{R\&}$$

$$A, \gamma_1, \ldots, \gamma_n \vdash_B \theta_1, \ldots, \theta_m \quad \text{and} \quad B, \gamma_1, \ldots, \gamma_n \vdash_B \theta_1, \ldots, \theta_m \tag{L$\vee$}$$

I'll here simply write $[\![\&\Gamma]\!]^P$ and $[\![\bigvee\Theta]\!]^C$ for brevity. Our inductive hypothesis guarantees the following:

$$(([\![\&\Gamma]\!]^P)^\curlyvee \sqcup ([\![\bigvee\Theta]\!]^C)^\curlyvee \sqcup ([\![A]\!]^P)^\curlyvee)^{\curlyvee\curlyvee} \subseteq \mathbb{I}_{\mathcal{M}} \quad \text{and} \quad (([\![\&\Gamma]\!]^P)^\curlyvee \sqcup ([\![\bigvee\Theta]\!]^C)^\curlyvee \sqcup ([\![B]\!]^C)^\curlyvee)^{\curlyvee\curlyvee} \subseteq \mathbb{I}_{\mathcal{M}} \tag{L$\to$}$$

$$(([\![\&\Gamma]\!]^P)^\curlyvee \sqcup ([\![\bigvee\Theta]\!]^C)^\curlyvee \sqcup ([\![A]\!]^C)^\curlyvee)^{\curlyvee\curlyvee} \subseteq \mathbb{I}_{\mathcal{M}} \quad \text{and} \quad (([\![\&\Gamma]\!]^P)^\curlyvee \sqcup ([\![\bigvee\Theta]\!]^C)^\curlyvee \sqcup ([\![B]\!]^C)^\curlyvee)^{\curlyvee\curlyvee} \subseteq \mathbb{I}_{\mathcal{M}} \tag{R\&}$$

$$(([\![\&\Gamma]\!]^P)^\curlyvee \sqcup ([\![\bigvee\Theta]\!]^C)^\curlyvee \sqcup ([\![A]\!]^P)^\curlyvee)^{\curlyvee\curlyvee} \subseteq \mathbb{I}_{\mathcal{M}} \quad \text{and} \quad (([\![\&\Gamma]\!]^P)^\curlyvee \sqcup ([\![\bigvee\Theta]\!]^C)^\curlyvee \sqcup ([\![B]\!]^P)^\curlyvee)^{\curlyvee\curlyvee} \subseteq \mathbb{I}_{\mathcal{M}} \tag{L$\vee$}$$

Next, recall that because '$\curlyvee\curlyvee$' functions as a closure operation, it follows that if $X \subseteq Y$ that $X^{\curlyvee\curlyvee} \subseteq Y^{\curlyvee\curlyvee}$. I also observe in general that $X \cap Y \subseteq X$ and $X \cap Y \subseteq Y$. Hence, we have the following:

$$[\Big(([\![\&\Gamma]\!]^P)^\curlyvee \sqcup ([\![\bigvee\Theta]\!]^C)^\curlyvee \sqcup ([\![A]\!]^P)^\curlyvee\Big) \cap \Big(([\![\&\Gamma]\!]^P)^\curlyvee \sqcup ([\![\bigvee\Theta]\!]^C)^\curlyvee \sqcup ([\![B]\!]^C)^\curlyvee\Big)]^{\curlyvee\curlyvee} \subseteq \mathbb{I}_{\mathcal{M}} \tag{L$\to$}$$

$$[\Big(([\![\&\Gamma]\!]^P)^\curlyvee \sqcup ([\![\bigvee\Theta]\!]^C)^\curlyvee \sqcup ([\![A]\!]^C)^\curlyvee\Big) \cap \Big(([\![\&\Gamma]\!]^P)^\curlyvee \sqcup ([\![\bigvee\Theta]\!]^C)^\curlyvee \sqcup ([\![B]\!]^C)^\curlyvee\Big)]^{\curlyvee\curlyvee} \subseteq \mathbb{I}_{\mathcal{M}} \tag{R\&}$$

$$[\Big(([\![\&\Gamma]\!]^P)^\curlyvee \sqcup ([\![\bigvee\Theta]\!]^C)^\curlyvee \sqcup ([\![A]\!]^P)^\curlyvee\Big) \cap \Big(([\![\&\Gamma]\!]^P)^\curlyvee \sqcup ([\![\bigvee\Theta]\!]^C)^\curlyvee \sqcup ([\![B]\!]^P)^\curlyvee\Big)]^{\curlyvee\curlyvee} \subseteq \mathbb{I}_{\mathcal{M}} \tag{L$\vee$}$$

Via Lemma A.2.16 we may rewrite the above as:

$$(([\![\&\Gamma]\!]^P)^\curlyvee \sqcup ([\![\bigvee\Theta]\!]^C)^\curlyvee \sqcup ([\![A]\!]^P \cap [\![B]\!]^C)^\curlyvee)^{\curlyvee\curlyvee} \subseteq \mathbb{I}_{\mathcal{M}} \tag{L$\to$}$$

$$(([\![\&\Gamma]\!]^P)^\curlyvee \sqcup ([\![\bigvee\Theta]\!]^C)^\curlyvee \sqcup ([\![A]\!]^C \cap [\![B]\!]^C)^\curlyvee)^{\curlyvee\curlyvee} \subseteq \mathbb{I}_{\mathcal{M}} \tag{R\&}$$

$$(([\![\&\Gamma]\!]^P)^\curlyvee \sqcup ([\![\bigvee\Theta]\!]^C)^\curlyvee \sqcup ([\![A]\!]^P \cap [\![B]\!]^P)^\curlyvee)^{\curlyvee\curlyvee} \subseteq \mathbb{I}_{\mathcal{M}} \tag{L$\vee$}$$

Next I note the following equivalences:

$$[\![A]\!]^P \cap [\![B]\!]^C = [\![B \to A]\!]^P$$
$$[\![A]\!]^C \cap [\![B]\!]^C = [\![A\&B]\!]^C$$
$$[\![A]\!]^P \cap [\![B]\!]^P = [\![A \vee B]\!]^P$$

Thus, it follows immediately that

$$B \to A, \gamma_1, \ldots, \gamma_n \vDash_B \theta_1, \ldots, \theta_m$$
$$\gamma_1, \ldots, \gamma_n \vDash_B \theta_1, \ldots, \theta_m, A\&B$$
$$A \vee B, \gamma_1, \ldots, \gamma_n \vDash_B \theta_1, \ldots, \theta_m$$

It therefore follows in all cases that:

$$\Gamma \vdash_B \Theta \Rightarrow \Gamma \vDash_B \Theta.$$

■

Theorem A.2.23 (Completeness)**.** *The sequent calculus is complete:*

$$\Gamma \vDash_B \Theta \Rightarrow \Gamma \vdash_B \Theta.$$

Proof. I show the contrapositive by constructing canonical models on which $\Gamma \vdash \Theta$ iff $\Gamma \vDash_{\mathcal{M}} \Theta$. Hence I wish to show that:

$$\Gamma \nvdash_B \Theta \Rightarrow \Gamma \nvDash_B \Theta.$$

Thus suppose that $\Gamma \nvdash_B \Theta$. I construct a model $\mathcal{M} = \langle \mathbf{P}, \mathbb{I}, [\![\cdot]\!] \rangle$ fit for B on which $\Gamma \nvDash_{\mathcal{M}} \Theta$.

Setting up the canonical models is fairly straightforward. We simply let $\mathbb{I}_{\mathcal{M}} = \vdash_B$, i.e.:[14]

$$\mathbb{I}_{\mathcal{M}} = \{\langle \Gamma, \Delta \rangle | \Gamma \vdash_B \Theta\}.$$

[14] I have chosen this because I consider it a more elegant solution. It would be equally sufficient, however, simply to stipulate that $\mathbb{I}_{\mathcal{M}} = \mid\!\sim_B \cap \mathcal{P}(\mathcal{L}_0)$ (i.e. contains all and only the atomic sequents). The crucial difference, however, concerns setting up the interpretation function.

Next, I prove that the interpretation function meets the following constraint:

$$[\![A]\!] = \langle \langle \{A\}, \emptyset \rangle^{\curlyvee}, \langle \emptyset, \{A\} \rangle^{\curlyvee} \rangle.$$

I previously defined interpretation functions such that the above always hold of *atomic* sentences, but it need not in general hold for all sentences (in all models). Really, what I'll be proving then is that in canonical models, the interpretation function works as follows:

$$[\![A]\!] = \langle \{\langle \Gamma, \Theta \rangle | A, \Gamma \vdash_B \Theta\}, \{\langle \Gamma, \Theta \rangle | \Gamma \vdash_B \Theta, A\} \rangle$$

Note that it should be immediately obvious that these two equations are equivalent given the choice of $\mathbb{I}_{\mathcal{M}}$. Once I've shown this it will follow that all and only good implications turn out to be semantic entailments.

In order to keep things straight, I will therefore introduce new notation for our intended interpretation function: $(\!|\cdot|\!)$. Thus let

$$(\!|A|\!) = \langle \{\langle \Gamma, \Theta \rangle | A, \Gamma \vdash_B \Theta\}, \{\langle \Gamma, \Theta \rangle | \Gamma \vdash_B \Theta, A\} \rangle = \langle \langle \{A\}, \emptyset \rangle^{\curlyvee}, \langle \emptyset, \{A\} \rangle^{\curlyvee} \rangle.$$

I now show that $(\!|A|\!) = [\![A]\!]$ by induction on logical complexity.

For the base case suppose A is atomic, then we have our result. Since $\mathbb{I}_{\mathcal{M}} = \vdash_B$ it follows that

$$\begin{aligned}
[\![A]\!] &= \langle \langle \{A\}, \emptyset \rangle^{\curlyvee}, \langle \emptyset, \{A\} \rangle^{\curlyvee} \rangle \\
&= \langle \{\langle \Gamma, \Theta \rangle | A, \Gamma \vdash_B \Theta\}, \{\langle \Gamma, \Theta \rangle | \Gamma \vdash_B \Theta, A\} \rangle \\
&= (\!|A|\!).
\end{aligned}$$

Thus suppose our result holds for sentences of logical complexity strictly less than n and now suppose that A has complexity of n. A is either a conditional, conjunction, disjunction, or negation. To make navigation easier I will use large headings for each of these cases. In each case I shall show $[\![A]\!] = (\!|A|\!)$ by showing containment in both directions for each member of the ordered pair. Note that $(\!|A|\!)^P$ and $(\!|A|\!)^C$ are used in the same way as $[\![A]\!]^P$ and $[\![A]\!]^C$.

Case 1: Conditional

Suppose A is a conditional, then it is of the form $B \to C$. I treat $[\![B \to C]\!]^P$ and $[\![B \to C]\!]^C$ separately. I am quite thorough in handling the conditional which can make following the "larger argument" more tedious. Consulting the corresponding proofs for the conjunction and disjunction may prove helpful.

I first show $[\![B \to C]\!]^P = (\!|B \to C|\!)^P$I. By our connective definition we have:

$$[\![B \to C]\!]^P = [\![B]\!]^C \cap [\![C]\!]^P.$$

By our inductive hypothesis, therefore:

$$[\![B]\!]^C \cap [\![C]\!]^P = \langle \emptyset, \{B\}\rangle^\curlyvee \cap \langle \{C\}, \emptyset\rangle^\curlyvee.$$

This is of course equivalent to

$$\langle \emptyset, \{B\}\rangle^\curlyvee \cap \langle \{C\}, \emptyset\rangle^\curlyvee = \{\langle \Gamma, \Theta\rangle | \Gamma \vdash_B \Theta, B \quad \text{and} \quad C, \Gamma \vdash_B \Theta\}.$$

In the sequent calculus our rule for L$\to$ has it that $\Gamma \vdash_B \Theta, B$ and $C, \Gamma \vdash_B \Theta$ iff $B \to C, \Gamma \vdash_B \Theta$. Thus we may rewrite the above:

$$\{\langle \Gamma, \Theta\rangle | \Gamma \vdash_B \Theta, B \quad \text{and} \quad C, \Gamma \vdash_B \Theta\} = \{\langle \Gamma, \Theta\rangle | B \to C, \Gamma \vdash_B \Theta\} \quad = \langle \{B \to C\}, \emptyset\rangle^\curlyvee = (\!|B \to C|\!)^P.$$

Thus

$$[\![B \to C]\!]^P = (\!|B \to C|\!)^P.$$

Next I show $[\![B \to C]\!]^C = (\!|B \to C|\!)^C$. From our semantic definition, we have

$$[\![B \to C]\!]^C = (([\![B]\!]^P)^\curlyvee \sqcup ([\![C]\!]^C)^\curlyvee)^\curlyvee.$$

From our inductive hypothesis, this is therefore

$$(([\![B]\!]^P)^\curlyvee \sqcup ([\![C]\!]^C)^\curlyvee)^\curlyvee = (\langle \{B\}, \emptyset\rangle^{\curlyvee\curlyvee} \sqcup \langle \emptyset, \{C\}\rangle^{\curlyvee\curlyvee})^\curlyvee.$$

($\subseteq$) Now, obviously $\langle \{B\}, \emptyset\rangle \in \langle \{B\}, \emptyset\rangle^{\curlyvee\curlyvee}$ and $C \in \langle \emptyset, \{C\}\rangle^{\curlyvee\curlyvee}$. Thus $\langle \{B\}, \{C\}\rangle \in \langle \{B\}, \emptyset\rangle^{\curlyvee\curlyvee} \sqcup \langle \emptyset, \{C\}\rangle^{\curlyvee\curlyvee}$.

Now, let $\langle \Delta, \Lambda\rangle \in [\![B \to C]\!]^C$ be arbitrary. Clearly $\langle \Delta, \Lambda\rangle \sqcup \langle \{B\}, \{C\}\rangle \in \mathbb{I}_\mathcal{M}$. Thus, $\langle \Delta, \Lambda\rangle \in \{\langle \Gamma, \Theta\rangle | \Gamma, B \vdash_B C, \Theta\}$. In the sequent calculus our rule for R$\to$ has it that if

$\Gamma, B \vdash_B C, \Theta$, then $\Gamma \vdash_B \Theta, A \to B$. Thus, $\langle \Delta, \Lambda \rangle \in \{\langle \Gamma, \Theta \rangle | \Gamma \vdash_B \Theta, B \to C\}$, i.e. $\langle \Delta, \Lambda \rangle \in \langle \emptyset, \{B \to C\} \rangle^{\curlyvee}$ and so in $(\!|B \to C|\!)^C$. It follows that

$$[\![B \to C]\!]^C \subseteq (\!|B \to C|\!)^C.$$

($\supseteq$) First we note that:

$$\begin{aligned} (\!|B \to C|\!)^C &= \langle \emptyset, \{B \to C\} \rangle^{\curlyvee} \\ &= \{\langle \Gamma, \Theta \rangle | \Gamma \vdash_B \Theta, B \to C\}. \end{aligned}$$

Since our rule for R$\to$, we have $\Gamma \vdash_B \Theta, B \to C$ iff $\Gamma, B \vdash_B C, \Theta$, thus

$$\begin{aligned} \{\langle \Gamma, \Theta \rangle | \Gamma \vdash_B \Theta, B \to C\} &= \{\langle \Gamma, \Theta \rangle | \Gamma, B \vdash_B C, \Theta\} \\ &= \langle \{B\}, \{C\} \rangle^{\curlyvee} \\ &= \langle \{B\}, \{C\} \rangle^{\curlyvee\curlyvee\curlyvee} \\ &= (\langle \{B\}, \emptyset \rangle \sqcup \langle \emptyset, \{C\} \rangle)^{\curlyvee\curlyvee\curlyvee}. \end{aligned}$$

Now, we have that via Lemma A.2.14 that $\langle \{B\}, \emptyset \rangle^{\curlyvee\curlyvee} \sqcup \langle \emptyset, \{C\} \rangle^{\curlyvee\curlyvee} \subseteq (\langle \{B\}, \emptyset \rangle \sqcup \langle \emptyset, \{C\} \rangle)^{\curlyvee\curlyvee}$. Hence via Corollary A.2.7 we have that

$$(\langle \{B\}, \emptyset \rangle \sqcup \langle \emptyset, \{C\} \rangle)^{\curlyvee\curlyvee\curlyvee} \subseteq (\langle \{B\}, \emptyset \rangle^{\curlyvee\curlyvee} \sqcup \langle \emptyset, \{C\} \rangle^{\curlyvee\curlyvee})^{\curlyvee}.$$

From our inductive hypothesis we have

$$\begin{aligned} (\langle \{B\}, \emptyset \rangle^{\curlyvee\curlyvee} \sqcup \langle \emptyset, \{C\} \rangle^{\curlyvee\curlyvee})^{\curlyvee} &= (([\![B]\!]^P)^{\curlyvee} \sqcup ([\![C]\!]^C)^{\curlyvee})^{\curlyvee} \\ &= [\![B \to C]\!]^C. \end{aligned}$$

Thus $(\!|B \to C|\!)^C \subseteq [\![B \to C]\!]^C$.

Since we have containment in both directions we therefore have $[\![B \to C]\!]^C = (\!|B \to C|\!)^C$ and thus

$$[\![B \to C]\!] = (\!|B \to C|\!).$$

Case 2: Conjunction

Suppose A is a conjunction, then it is of the form $B\&C$. I treat $[\![B\&C]\!]^P$ and $[\![B\&C]\!]^C$ separately. Since both cases are analogous to the proofs concerning the conditional I will be much briefer here.

To start we wish to show $[\![B\&C]\!]^P = (\!|B\&C|\!)^P$. The proof is similar to establishing the analogous result for the conclusory role of the conditional. In that case, the proof can be reduced to establishing one key step, which I mark here using (***):

$$\begin{aligned}
[\![B\&C]\!]^P &= (([\![B]\!]^P)^\curlyvee \sqcup ([\![C]\!]^P)^\curlyvee)^\curlyvee \\
&= (\langle\{B\},\emptyset\rangle^{\curlyvee\curlyvee} \sqcup \langle\{C\},\emptyset\rangle^{\curlyvee\curlyvee})^\curlyvee && \text{(Inductive Hypothesis)} \\
&\stackrel{?}{=} (\langle\{B\},\emptyset\rangle \sqcup \langle\{C\},\emptyset\rangle)^{\curlyvee\curlyvee\curlyvee} && (***) \\
&= \langle\{B,C\},\emptyset\rangle^{\curlyvee\curlyvee\curlyvee} \\
&= \langle\{B,C\},\emptyset\rangle^{\curlyvee} \\
&= \{\langle\Gamma,\Theta\rangle | A, B, \Gamma \vdash_B \Theta\} \\
&= \{\langle\Gamma,\Theta\rangle | A\&B, \Gamma \vdash_B \Theta\} && \text{(L\& Rule)} \\
&= \langle\{B\&C\},\emptyset\rangle^{\curlyvee} \\
&= (\!|B\&C|\!)^P.
\end{aligned}$$

Thus the proof of $[\![B\&C]\!]^P = (\!|B\&C|\!)^P$ reduces to

$$(\langle\{B\},\emptyset\rangle^{\curlyvee\curlyvee} \sqcup \langle\{C\},\emptyset\rangle^{\curlyvee\curlyvee})^\curlyvee \stackrel{?}{=} (\langle\{B\},\emptyset\rangle \sqcup \langle\{C\},\emptyset\rangle)^{\curlyvee\curlyvee\curlyvee}. \qquad (***)$$

($\subseteq$) Let $\langle\Delta,\Lambda\rangle \in (\langle\{B\},\emptyset\rangle^{\curlyvee\curlyvee} \sqcup \langle\{C\},\emptyset\rangle^{\curlyvee\curlyvee})^\curlyvee$ be arbitrary. Next, observe that $\langle\{B\},\emptyset\rangle \in \langle\{B\},\emptyset\rangle^{\curlyvee\curlyvee}$ and $\langle\{C\},\emptyset\rangle \in \langle\{C\},\emptyset\rangle^{\curlyvee\curlyvee}$ and hence $\langle\{B\},\emptyset\rangle \sqcup \langle\{C\},\emptyset\rangle \in \langle\{B\},\emptyset\rangle^{\curlyvee\curlyvee} \sqcup \langle\{C\},\emptyset\rangle^{\curlyvee\curlyvee}$. It follows that $\langle\{B\},\emptyset\rangle \sqcup \langle\{C\},\emptyset\rangle \sqcup \langle\Delta,\Lambda\rangle \in \mathbb{I}_{\mathcal{M}}$ and thus $\langle\Delta,\Lambda\rangle \in (\langle\{B\},\emptyset\rangle \sqcup \langle\{C\},\emptyset\rangle)^\curlyvee$.

($\supseteq$) We have (via Lemma A.2.14) that

$$\langle\{B\},\emptyset\rangle^{\curlyvee\curlyvee} \sqcup \langle\{C\},\emptyset\rangle^{\curlyvee\curlyvee} \subseteq (\langle\{B\},\emptyset\rangle \sqcup \langle\{C\},\emptyset\rangle)^{\curlyvee\curlyvee}$$

Hence via Corollary A.2.7 we have:

$$(\langle\{B\},\emptyset\rangle \sqcup \langle\{C\},\emptyset\rangle)^{\curlyvee\curlyvee\curlyvee} \subseteq (\langle\{B\},\emptyset\rangle^{\curlyvee\curlyvee} \sqcup \langle\{C\},\emptyset\rangle^{\curlyvee\curlyvee})^\curlyvee.$$

Hence,

$$[\![B\&C]\!]^P = (\!|B\&C|\!)^P.$$

Next, I show that $[\![B\&C]\!]^C = (\!|B\&C|\!)^C$. Since this case is very similar to the analogous result for the premissory role of the conditional, I allow myself to be much briefer here. I reason as follows:

$$\begin{aligned}
[\![B\&C]\!]^C &= [\![B]\!]^C \cap [\![C]\!]^C \\
&= \langle \emptyset, \{B\}\rangle^{\curlyvee} \cap \langle \emptyset, \{C\}\rangle^{\curlyvee} && \text{(Inductive Hypothesis)} \\
&= \{\langle \Gamma, \Theta\rangle | \Gamma \vdash_B \Theta, B \quad \text{and} \quad \Gamma \vdash_B \Theta, C\} \\
&= \{\langle \Gamma, \Theta\rangle | \Gamma \vdash_B \Theta, B\&C\} && \text{(R\&)} \\
&= \langle \emptyset, \{B\&C\}\rangle^{\curlyvee} && = (\!|B\&C|\!)^C.
\end{aligned}$$

We therefore have that $[\![B\&C]\!]^C = (\!|B\&C|\!)^C$, and thus:

$$[\![B\&C]\!] = (\!|B\&C|\!).$$

Case 3: Disjunction

Suppose A is a disjunction, then it is of the form $B \vee C$. I treat $[\![B \vee C]\!]^P$ and $[\![B \vee C]\!]^C$ separately. Since both cases are analogous to the proofs concerning the conditional and conjunction I will be much briefer here.

I first show that $[\![B \vee C]\!]^P = (\!|B \vee C|\!)^P$. Since this case is very similar to the analogous result for the premissory role of the conditional and conclusory role of the conjunction, I allow myself to be much briefer here. I reason as follows:

$$\begin{aligned}
[\![B \vee C]\!]^P &= [\![B]\!]^P \cap [\![C]\!]^P \\
&= \langle \{B\}, \emptyset\rangle^{\curlyvee} \cap \langle \{C\}, \emptyset\rangle^{\curlyvee} && \text{(Inductive Hypothesis)} \\
&= \{\langle \Gamma, \Theta\rangle | B, \Gamma \vdash_B \Theta \quad \text{and} \quad C, \Gamma \vdash_B \Theta\} \\
&= \{\langle \Gamma, \Theta\rangle | B \vee C, \Gamma \vdash_B \Theta\} && \text{(L}\vee\text{)} \\
&= \langle \{B \vee C\}, \emptyset\rangle^{\curlyvee} && = (\!|B \vee C|\!)^P.
\end{aligned}$$

We therefore have that

$$[\![B \vee C]\!]^P = (\!|B \vee C|\!)^P.$$

Next I show $[\![B \vee C]\!]^C = (\!|B \vee C|\!)^C$. I allow myself to be briefer here. As with the conjunction, we can mark the crucial step with (***):

$$\begin{aligned}
[\![B \vee C]\!]^C &= (([\![B]\!]^C)^\curlyvee \sqcup ([\![C]\!]^C)^\curlyvee)^\curlyvee \\
&= (\langle \emptyset, \{B\} \rangle^{\curlyvee\curlyvee} \sqcup \langle \emptyset, \{C\} \rangle^{\curlyvee\curlyvee})^\curlyvee && \text{(Inductive Hypothesis)} \\
&\stackrel{?}{=} (\langle \emptyset, \{B\} \rangle \sqcup \langle \emptyset, \{C\} \rangle)^{\curlyvee\curlyvee\curlyvee} && \text{(***)} \\
&= \langle \emptyset, \{B, C\} \rangle^{\curlyvee\curlyvee\curlyvee} \\
&= \langle \emptyset, \{B, C\} \rangle^{\curlyvee} \\
&= \{\langle \Gamma, \Theta \rangle | \Gamma \vdash_B \Theta, B, C\} \\
&= \{\langle \Gamma, \Theta \rangle | \Gamma \vdash_B \Theta, B, C\} && \text{(R}\vee\text{ Rule)} \\
&= \langle \emptyset, \{B \vee C\} \rangle^{\curlyvee} \\
&= (\!|B \vee C|\!)^C.
\end{aligned}$$

Thus the proof reduces to

$$(\langle \emptyset, \{B\} \rangle^{\curlyvee\curlyvee} \sqcup \langle \emptyset, \{C\} \rangle^{\curlyvee\curlyvee})^\curlyvee \stackrel{?}{=} (\langle \emptyset, \{B\} \rangle \sqcup \langle \emptyset, \{C\} \rangle)^{\curlyvee\curlyvee\curlyvee}.$$

($\subseteq$) Let $\langle \Delta, \Lambda \rangle \in (\langle \emptyset, \{B\} \rangle^{\curlyvee\curlyvee} \sqcup \langle \emptyset, \{C\} \rangle^{\curlyvee\curlyvee})^\curlyvee$. Clearly $\langle \emptyset, \{B\} \rangle \sqcup \langle \emptyset, \{C\} \rangle \in \langle \emptyset, \{B\} \rangle^{\curlyvee\curlyvee} \sqcup \langle \emptyset, \{C\} \rangle^{\curlyvee\curlyvee}$. Thus $\langle \emptyset, \{B\} \rangle \sqcup \langle \emptyset, \{C\} \rangle \sqcup \langle \Delta, \Lambda \rangle \in \mathbb{I}_{\mathcal{M}}$.

($\supseteq$) We have (via Lemma A.2.14)

$$(\langle \emptyset, \{B\} \rangle^{\curlyvee\curlyvee} \sqcup \langle \emptyset, \{C\} \rangle^{\curlyvee\curlyvee}) \subseteq (\langle \emptyset, \{B\} \rangle \sqcup \langle \emptyset, \{C\} \rangle)^{\curlyvee\curlyvee},$$

and thus via Corollary A.2.7:

$$(\langle \emptyset, \{B\} \rangle \sqcup \langle \emptyset, \{C\} \rangle)^{\curlyvee\curlyvee\curlyvee} \subseteq (\langle \emptyset, \{B\} \rangle^{\curlyvee\curlyvee} \sqcup \langle \emptyset, \{C\} \rangle^{\curlyvee\curlyvee})^\curlyvee.$$

Hence, $[\![B \vee C]\!]^C = (\!|B \vee C|\!)^C$, and thus

$$[\![B \vee C]\!] = (\!|B \vee C|\!).$$

Case 4: Negation

Finally, suppose that A is a negated sentence, i.e. A has the form $\neg B$. We must show $[\![\neg B]\!] = (\!|\neg B|\!)$. I here treat the cases of $[\![B]\!]^P$ and $[\![B]\!]^C$ simultaneously since they are handled analogously and since the general principles of the proof should be familiar from the previous cases. I reason as follows:

$$\begin{aligned}
[\![\neg B]\!] &= \langle [\![B]\!]^C, [\![B]\!]^P \rangle \\
&= \langle \langle \emptyset, \{B\}\rangle^\curlyvee, \langle \{B\}, \emptyset\rangle^\curlyvee \rangle && \text{(Inductive Hypothesis)} \\
&= \langle \{\langle \Gamma, \Theta\rangle | \Gamma \vdash_B \Theta, B\}, \{\langle \Gamma, \Theta\rangle | B, \Gamma \vdash_B \Theta\}\rangle \\
&= \langle \{\langle \Gamma, \Theta\rangle | \neg B, \Gamma \vdash_B \Theta\}, \{\langle \Gamma, \Theta\rangle | \Gamma \vdash_B \Theta, \neg B\}\rangle && \text{(L}\neg\text{/ R}\neg\text{)} \\
&= \langle \langle \{\neg B\}, \emptyset\rangle^\curlyvee, \langle \emptyset, \{\neg B\}\rangle^\curlyvee \rangle \\
&= (\!|B|\!).
\end{aligned}$$

Note that the step labeled (L$\neg$/R$\neg$) follows since in our sequent calculus $\Gamma \vdash_B \Theta, B$ iff $\neg B, \Gamma \vdash_B \Theta$ and likewise $B, \Gamma \vdash_B \Theta$ iff $\Gamma \vdash_B \Theta, \neg B$, i.e. our '$\neg$' is involutive.

We therefore have that $[\![\neg B]\!] = (\!|\neg B|\!)$. Since we have shown this in all cases it follows that

$$[\![A]\!] = (\!|A|\!).$$

Main Proof Body

Recall the point of establishing this about our canonical models was to show that $\Gamma \vdash_B \Theta$ iff $\Gamma \vDash_{\mathcal{M}} \Theta$ on such models. I next show this. Since many of the principles of the proof are similar to what has already been shown I allow myself some brevity here. Let us enumerate Γ and Θ as $\gamma_1, \ldots, \gamma_n$ and $\theta_1, \ldots, \theta_m$. I reason as follows

$$\begin{aligned}
\Gamma \vDash_{\mathcal{M}} \Theta &\Leftrightarrow (([\![\gamma_1]\!]^P)^\curlyvee \sqcup \cdots \sqcup ([\![\gamma_n]\!]^P)^\curlyvee \sqcup ([\![\theta_1]\!]^C)^\curlyvee \sqcup \cdots \sqcup ([\![\theta_m]\!]^C)^\curlyvee)^{\curlyvee\curlyvee} \subseteq \mathcal{I}_{\mathcal{M}} \\
&\Leftrightarrow (((\!|\gamma_1|\!)^P)^\curlyvee \sqcup \cdots \sqcup ((\!|\gamma_n|\!)^P)^\curlyvee \sqcup ((\!|\theta_1|\!)^C)^\curlyvee \sqcup \cdots \sqcup ((\!|\theta_m|\!)^C)^\curlyvee)^{\curlyvee\curlyvee} \subseteq \mathcal{I}_{\mathcal{M}} \\
&\Leftrightarrow (\langle \{\gamma_1\}, \emptyset\rangle^{\curlyvee\curlyvee} \sqcup \cdots \sqcup \langle \{\gamma_n\}, \emptyset\rangle^{\curlyvee\curlyvee} \sqcup \langle \{\theta_1\}, \emptyset\rangle^{\curlyvee\curlyvee} \sqcup \cdots \sqcup \langle \{\theta_m\}, \emptyset\rangle^{\curlyvee\curlyvee})^{\curlyvee\curlyvee} \subseteq \mathcal{I}_{\mathcal{M}} \\
&\Leftrightarrow (\langle \{\gamma_1\}, \emptyset\rangle \sqcup \cdots \sqcup \langle \{\gamma_n\}, \emptyset\rangle \sqcup \langle \{\theta_1\}, \emptyset\rangle \sqcup \cdots \sqcup \langle \{\theta_m\}, \emptyset\rangle)^{\curlyvee\curlyvee\curlyvee\curlyvee} \subseteq \mathcal{I}_{\mathcal{M}} \\
&\Leftrightarrow \langle \{\gamma_1, \ldots, \gamma_n\}, \{\theta_1, \ldots, \theta_m\}\rangle^{\curlyvee\curlyvee} \subseteq \mathcal{I}_{\mathcal{M}} \\
&\Leftrightarrow \langle \Gamma, \Theta\rangle^{\curlyvee\curlyvee} \subseteq \mathcal{I}_{\mathcal{M}} \\
&\Leftrightarrow \langle \Gamma, \Theta\rangle \in \mathbb{I}_{\mathcal{M}} && \text{(A.2.8)} \\
&\Leftrightarrow \Gamma \vdash_B \Theta.
\end{aligned}$$

It follows that

$$\Gamma \vDash_{\mathcal{M}} \Theta \Leftrightarrow \Gamma \vdash_B \Theta.$$

Thus if $\Gamma \nvdash_B \Theta$, then $\Gamma \nvDash_B \Theta$. Hence

$$\Gamma \vDash_B \Theta \Rightarrow \Gamma \vdash_B \Theta.$$

■

From these two theorems we have

$$\Gamma \vdash_B \Theta \Leftrightarrow \Gamma \vDash_B \Theta,$$

that is, the sequent calculus is sound and complete with respect to the inferential role semantics.

A.2.4 Functional Completeness of Connectives

In this section I want to distinguish two senses in which the connectives may be "functionally complete". The first sense is meant to pick out the idea that for any $X^{\curlyvee\curlyvee} \subseteq \mathbf{P}$, we may find a sentence $A \in \mathcal{L}$ such that either $[\![A]\!]^P = X^{\curlyvee\curlyvee}$ or $[\![A]\!]^C = X^{\curlyvee\curlyvee}$. In other words, the connectives are expressive enough that any closed set will play the role of either premises or conclusion for *some* sentence.

The second sense has it that for any $\langle X^{\curlyvee\curlyvee}, Y^{\curlyvee\curlyvee}\rangle \subseteq \mathbf{P}^2$ (i.e. for any *proper inferential role*) there is some sentence that picks out *exactly that role.*

I define the notions of functional completeness. I simply call them "functional completeness 1" (Φ_1) and "functional completeness 2" (Φ_2) for now (lacking a better name).

Some important things to note:

- Functional completeness is a property that obtains between a base consequence relation and a set of connectives. Thus, my connectives $\{\&, \vee, \rightarrow, \neg\}$ are functionally complete in neither sense, but if we restrict B to classical logic: minimal, flat, then $\{\&, \vee, \rightarrow, \neg\}$ *is* functionally complete in the first, but not the second sense.

- This notion will be important in the section following this since Φ_1 turns out to be sufficient for $\Delta_4 \Rightarrow \Delta_3$. Φ_2 is important since it turns out that Φ_2 is sufficient for $\Delta_3 \Rightarrow \Delta_2$.
- Φ_2 isn't necessarily a desirable quality.

A.2.5 Theories

In line with material pursued in the previous appendix, we can also limit models to realize or satisfy theories. I have already detailed in the proof theory how this works and since the model theory is complete we can obviously perform the same sorts of restrictions within it (i.e. we can simply impose constraints on $\vdash_0$, e.g. on models fit for $B = \vdash_0$).

In the case where we impose theories (i.e. sets of sentences of inferential relations) how this will work is obvious. We simply read off from the theory/inferential relation what must be included in B (or what B must be) to ensure that we get the proper restriction on models.

Analogously we can read similar *structural* constraints in the same way. Since the structural constraints I considered in the previous section were simply reducible to constraints on B, we can make similar enrichments here.

Bibliography

Aarnio, M. L. (2010). Unreasonable knowledge. *Philosophical Perspectives*, 24:1–21.

Alchourrón, C. E., Gärdenfors, P., and Makinson, D. (1985). On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50(2):510–530.

Allais, M. (1953). Le comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'école américaine. *Econometrica: Journal of the Econometric Society*, pages 503–546.

Andreoli, J.-M. (1992). Logic programming with focusing proofs in linear logic. *Journal of Logic and Computation*, 2(3):297–347.

Antonelli, G. A. (1999). A directly cautious theory of defeasible consequence for default logic via the notion of general extension. *Artificial Intelligence*, 109(1-2):71–109.

Arieli, O. and Avron, A. (2000). General patterns for nonmonotonic reasoning: From basic entailments to plausible relations. *Logic Journal of the IGPL*, 8(2):119–148.

Bach, K. (2005). Context ex machina. In Szabó, Z. G., editor, *Semantics versus pragmatics*. Oxford University Press Oxford.

Baker-Hytch, M. and Benton, M. A. (2015). Defeatism defeated. *Philosophical Perspectives*, 29(1):40–66.

Barker, C. (2010). Free choice permission as resource-sensitive reasoning. *Semantics and Pragmatics*, 3:10–1.

Barrio, E., Rosenblatt, L., and Tajer, D. (2015a). The logics of strict-tolerant logic. *Journal of Philosophical Logic*, 44(5):551–571.

Barrio, E., Rosenblatt, L., and Tajer, D. (2015b). The logics of strict-tolerant logic. *Journal of Philosophical Logic*, 44(5):551–571.

Barrio, E., Rosenblatt, L., and Tajer, D. (2016). Capturing naive validity in the cut-free approach. *Synthese*, pages 1–17.

Barwise, J. and Etchemendy, J. (1987). *The liar: An essay on truth and circularity*. Oxford University Press.

Batens, D. (2004). The need for adaptative logics in epistemology. In Rahman, S., Symons, J., Gabbay, D., and Bendegem, J., editors, *Logic, Epistemology, and the Unity of Science*, pages 459–485. Kluwer Academic Publishers.

Beall, J. (2007). Truth and paradox: a philosophical sketch. In *Philosophy of logic*, pages 325–410. Elsevier.

Beall, J. (2009). *Spandrels of truth*. Oxford University Press.

Beall, J. and Murzi, J. (2013). Two flavors of curry's paradox. *The Journal of Philosophy*, 110(3):143–165.

Benton, M. A. (2016). Knowledge and evidence you should have had. *Episteme*, 13(4):471–479.

Bergmann, M. (1997). *Internalism, Externalism and Epistemic Defeat*. PhD thesis.

Bergmann, M. (2006). *Justification Without Awareness: A Defense of Epistemic Externalism*. Oxford University Press.

Bezuidenhout, A. (2009). Contextualism and the role of contextual frames. *Manuscrito: Revista Internacional de Filosofia*, 32(1):59–84.

Bezuidenhout, A. (2017). Contextualism and semantic minimalism. In Huang, Y., editor, *The Oxford handbook of pragmatics*, pages 21–46. Oxford University Press.

Bimbo, K. (2006). Curry-type paradoxes. *Logique et Analyse*, pages 227–240.

Borg, E. (2004). *Minimal semantics*. Oxford University Press.

Borg, E. (2007). Minimalism versus contextualism in semantics. In Preyer, G. and Peter, G., editors, *Context-sensitivity and semantic minimalism: New essays on semantics and pragmatics*, volume 546, pages 339–360. Oxford: Oxford University Press.

Brady, R. (2006). *Universal Logic*. CSLI Publications.

Brandom, R. (1983). Asserting. *Noûs*, 17(4):637–50.

Brandom, R. (1998). *Making it explicit: Reasoning, representing, and discursive commitment*. Harvard university press.

Brandom, R. (2008). *Between saying and doing: Towards an analytic pragmatism*. Oxford University Press.

Brandom, R. (2009). *Articulating reasons*. Harvard University Press.

Brandom, R. (2018a). From logical expressivism to expressivist logic: Sketch of a program and some implementations. *Philosophical Issues*, 28(1):70–88.

Brandom, R. (2018b). From logical expressivism to expressivist logics: Sketch of a program and some implementations. In *From Rules to Meanings*, pages 151–164. Routledge.

Braun, D. (2017). Indexicals. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Summer 2017 edition.

Broome, J. (1991). *Weighing Goods: Equality, Uncertainty and Time*. Wiley-Blackwell.

Burge, T. (1979). Semantical paradox. *The Journal of Philosophy*, 76(4):169–198.

Cappelen, H. and Lepore, E. (2008). *Insensitive semantics: A defense of semantic minimalism and speech act pluralism*. John Wiley & Sons.

Caret, C. R. and Weber, Z. (2015). A note on contraction-free logic for validity. *Topoi*, 34(1):63–74.

Casullo, A. (2018). Pollock and sturgeon on defeaters. *Synthese*, 195(7):2897–2906.

Chisholm, R. M. (1964). The ethics of requirement. *American Philosophical Quarterly*, 1(2):147–153.

Chisholm, R. M. (1966). *Theory of Knowledge*.

Cobreros, P., Egré, P., Ripley, D., and van Rooij, R. (2012a). Tolerance and mixed consequence in the s'valuationist setting. *Studia logica*, 100(4):855–877.

Cobreros, P., Egré, P., Ripley, D., and van Rooij, R. (2012b). Tolerant, classical, strict. *Journal of Philosophical Logic*, 41(2):347–385.

Cobreros, P., Egré, P., Ripley, D., and van Rooij, R. (2017). Tolerant reasoning: Nontransitive or nonmonotonic? *Synthese*, pages 1–25.

Cobreros, P., Egré, P., Ripley, D., and van Rooij, R. (2020). Inferences and metainferences in st. *Journal of Philosophical Logic*, 49(6):1057–1077.

Curry, H. B. (1942). The inconsistency of certain formal logics. *The Journal of Symbolic Logic*, 7(3):115–117.

Da Ré, B. and Rosenblatt, L. (2018). Contraction, infinitary quantifiers, and omega paradoxes. *Journal of Philosophical Logic*, 47(4):611–629.

Dancy, J. (2004a). Enticing reasons. In Wallace, R. J., Pettit, P., Scheffler, S., and Smith, M., editors, *Reason and Value: Themes From the Moral Philosophy of Joseph Raz*, pages 91–118. Clarendon Press.

Dancy, J. (2004b). *Ethics without principles*. Oxford University Press on Demand.

DeRose, K. (1995). Solving the skeptical problem. *The philosophical review*, 104(1):1–52.

Dicher, B. (2016). A proof-theoretic defence of meaning-invariant logical pluralism. *Mind*, 125(499):727–757.

Dicher, B. (2020). Hopeful monsters: A note on multiple conclusions. *Erkenntnis*, 85(1):77–98.

Dummett, M. (1976). What is a theory of meaning? (ii). In Evans, G. and McDowell, J., editors, *Truth and Meaning: Essays in Semantics*. Oxford: Clarendon Press.

Dummett, M. A. E. (1975). What is a theory of meaning? In Guttenplan, S., editor, *Mind and Language*. Oxford University Press.

Field, H. (2008). *Saving truth from paradox*. Oxford University Press.

Frege, G. (1980). *The Foundations of Arithmetic: A Logico-Mathematical Enquiry into the Concept of Number*. Northwestern University Press, Evanston, IL, 2 edition.

Frege, G. (1988). *Die Grundlagen der Arithmetik*. Felix Meiner Verlag.

French, R. (2016). Structural reflexivity and the paradoxes of self-reference. *Ergo, an Open Access Journal of Philosophy*, 3.

Gabbay, D. M. (1985). Theoretical foundations for non-monotonic reasoning in expert systems. In *Logics and models of concurrent systems*, pages 439–457. Springer.

Gauker, C. (2012). What tipper is ready for: A semantics for incomplete predicates. *Noûs*, 46(1):61–85.

Geach, P. T. (1955). On insolubilia. *Analysis*, 15(3):71–72.

Geisst, C. R. (2013). *Beggar thy neighbor: A history of usury and debt*. University of Pennsylvania Press.

Gettier, E. L. (1963). Is justified true belief knowledge? *Analysis*, 23(6):121–123.

Gillon, B. S. (2015). Optional complements of english verbs and adjectives. In *Cognitive Science Perspectives on Verb Representation and Processing*, pages 67–75. Springer.

Girard, J.-Y. (1987). Linear logic. *Theoretical computer science*, 50(1):1–101.

Girard, J.-Y. (1995). Linear logic: its syntax and semantics. In *Advances in Linear Logic*, pages 1–42. Cambridge University Press.

Girard, J.-Y. (2011). *The Blind Spot: lectures on logic*. European Mathematical Society.

Glanzberg, M. (2001). The liar in context. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 103(3):217–251.

Glanzberg, M. (2004). A contextual-hierarchical approach to truth and the liar paradox. *Journal of Philosophical Logic*, 33(1):27–88.

Goldberg, S. (2017). Should have known. *Synthese*, 194(8):2863–2894.

Goldberg, S. C. (2016). On the epistemic significance of evidence you should have had. *Episteme*, 13(4):449–470.

Grice, H. P. (1975). Logic and conversation. In et. al., C., editor, *Syntax and semantics 3: Speech arts*. Elsevier.

Grice, P. (1957). Meaning. *Philosophical Review*, 66:377–388.

Grice, P. H. (1991). *Studies in the Way of Words*. Harvard University Press.

Grundmann, T. (2011). Defeasibility theory. In Bernecker, S. and Pritchard, D., editors, *The Routledge Companion to Epistemology*, pages 156–166. Routledge.

Gupta, A. (1982). Truth and paradox. *Journal of philosophical logic*, 11(1):1–60.

Gupta, A. and Belnap, N. D. (1993). *The revision theory of truth*. MIT Press.

Hansson, S. O. (2017). Logic of Belief Revision. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Winter 2017 edition.

Hart, H. (1948). The ascription of responsibility and rights. In *Proceedings of the Aristotelian society*, volume 49, pages 171–194.

Hlobil, U. (2016). A nonmonotonic sequent calculus for inferentialist expressivists. In Arazim, P. and Danák, M., editors, *The Logica Yearbook 2015*, pages 87–105. College Publications.

Horty, J. (2007a). Defaults with priorities. *Journal of Philosophical Logic*, 36(4):367–413.

Horty, J. F. (2001). Argument construction and reinstatement in logics for defeasible reasoning. *Artificial intelligence and Law*, 9(1):1–28.

Horty, J. F. (2007b). Reasons as defaults. *Philosopher's Imprint*, (7(3)).

Horty, J. F. (2012). *Reasons as defaults*. Oxford University Press.

Humberstone, I. L. (1988). Heterogeneous logic. *Erkenntnis*, 29(3):395–435.

Humberstone, L. (2006). Variations on a theme of curry. *Notre Dame Journal of Formal Logic*, 47(1):101–131.

Humberstone, L. (2007). Investigations into a left-structural right-substructural sequent calculus. *Journal of Logic, Language and Information*, 16(2):141–171.

Jeffrey, R. C. (1990). *The logic of decision*. University of Chicago press.

Kagan, S. (1989). *The Limits of Morality*. Oxford University Press.

Kant, I. (1900). Kants gesammelte schriften, akademie ausgabe, edited by the koeniglichen preussischen akademie der wissenschaften, 29 vols.

Kaplan, D. (1979). On the logic of demonstratives. *Journal of philosophical logic*, 8(1):81–98.

Kaplan, D. (2018). A multi-succedent sequent calculus for logical expressivists. In Arazim, P. and Lávička, T., editors, *The Logica Yearbook 2017*, pages 139–153. College Publications.

Kennedy, C. (2007). Vagueness and grammar: The semantics of relative and absolute gradable adjectives. *Linguistics and philosophy*, 30(1):1–45.

Kennedy, C. (2019). *Ambiguity and vagueness: An overview*, pages 236–271. De Gruyter Mouton.

Ketonen, O. (1944). *Untersuchungen zum prädikatenkalkü l.* Suomalainen tiedeakatemia.

Klein, P. D. (1971). A proposed definition of propositional knowledge. *the Journal of Philosophy*, 68(16):471–482.

Klein, P. D. (1976). Knowledge, causality, and defeasibility. *The Journal of Philosophy*, 73(20):792–812.

Koons, R. (2017). Defeasible reasoning. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, winter 2017 edition.

Kotzen, M. (2019). A formal account of epistemic defeat. In Braden, C., Borges, R., and Fitelson, B., editors, *Themes From Klein*. Springer Verlag.

Kratzer, A. (2012). *Modals and conditionals: New and revised perspectives*, volume 36. Oxford University Press.

Kraus, S., Lehmann, D., and Magidor, M. (1990). Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial intelligence*, 44(1-2):167–207.

Kripke, S. A. (1975). Outline of a theory of truth. *Journal of Philosophy*, 72(19):690–716.

Kvanvig, J. (2007). Two approaches to epistemic defeat. In Baker, D., editor, *Alvin Plantinga*, pages 107–124. Cambridge: Cambridge University Press.

Lackey, J. (2008). *Learning From Words: Testimony as a Source of Knowledge*. Oxford: Oxford University Press.

Lance, M. and Little, M. (2006). Particularism and antitheory. In Copp, D., editor, *The Oxford Handbook of Ethical Theory*, pages 567–594. Oxford University Press.

Lasonen-Aarnio, M. (2014). Higher-order evidence and the limits of defeat. *Philosophy and Phenomenological Research*, 88(2):314–345.

Lehrer, K. (1965). Knowledge, truth and evidence. *Analysis*, 25(5):168–175.

Lehrer, K. (1970). The fourth condition of knowledge: A defense. *The Review of Metaphysics*, 24(1):122–128.

Lehrer, K. (2019). *Exemplars of Truth*. Oxford University Press.

Lehrer, K. and Paxson Jr, T. (1969). Knowledge: Undefeated justified true belief. *The Journal of Philosophy*, 66(8):225–237.

Levy, S. R. (1977). Defeasibility theories of knowledge. *Canadian Journal of Philosophy*, 7(1):115–123.

Lewis, D. (1973). *Counterfactuals*. Harvard University Press.

Lewis, D. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic*, 8(1):339–59.

Lewis, D. (1982). Logic for equivocators. *Noûs*, 16(3):431–441.

Lewis, D. (1996). Elusive knowledge. *Australasian journal of Philosophy*, 74(4):549–567.

Lewis, D. K. (1970). General semantics. *Synthese*, 22(1-2):18–67.

Liang, C. and Miller, D. (2009). Focusing and polarization in linear, intuitionistic, and classical logics. *Theoretical Computer Science*, 410(46):4747–4768.

Löb, M. H. (1955). Solution of a problem of leon henkin. *The Journal of Symbolic Logic*, 20(2):115–118.

MacFarlane, J. (2011). What is assertion. In Brown, J. and Cappelen, H., editors, *Assertion*. Oxford University Press.

Makinson, D. (1994). General patterns in nonmonotonic reasoning. In Gabbay, H. and Robinson, editors, *Handbook of Logic in Artificial Intelligence and Logic Programming*, volume 3, pages 35–110. Oxford University Press.

Makinson, D. (2005). *Bridges from classical to nonmonotonic logic*. King's College Publications.

Malinowski, G. (2004). Inferential intensionality. *Studia Logica*, 76(1):3–16.

Malinowski, G. (2014). Kleene logic and inference. *Bulletin of the Section of Logic*, 43(1/2):43–52.

Mares, E. and Paoli, F. (2014). Logical consequence and the paradoxes. *Journal of Philosophical Logic*, 43(2-3):439–469.

McCarthy, J. M. and Hayes, P. J. (1969). Some philosophical problems from the standpoint of artificial intelligence. In Meltzer, B. and Mitchie, D., editors, *Machine Intelligence*, number 4. Edinburgh University Press.

Meadows, T. (2014). Fixed points for consequence relations. *Logique et Analyse*, pages 333–357.

Melis, G. (2014). Understanding undermining defeat. *Philosophical Studies*, 170(3):433–442.

Moretti, L. and Piazza, T. (2018). Defeaters in current epistemology: Introduction to the special issue. *Synthese*, 195(7):2845–2854.

Muñoz, D. (2019). Defeaters and disqualifiers. *Mind*, 128(511):887–906.

Murzi, J. and Rossi, L. (2021). Non-reflexivity and revenge. *Journal of Philosophical Logic*, pages 1–18.

Negri, S., Von Plato, J., and Ranta, A. (2008). *Structural proof theory*. Cambridge University Press.

Nicolai, C. and Rossi, L. (2018). Principles for object-linguistic consequence: From logical to irreflexive. *Journal of Philosophical Logic*, 47(3):549–577.

Partee, B. H. (1989). Binding implicit variables in quantified contexts. *University of Massachusetts Occasional Papers in Linguistics*, 15(2):12.

Peregrin, J. (2014). *Inferentialism: Why Rules Matter*. Palgrave Macmillan.

Plucknett, T. F. (1956). *A Concise History of the Common Law*. Little, Brown and Co.

Pollock, J. and Cruz, J. (1999). *Contemporary Theories of Knowledge*, chapter Externalism, pages 93–149. Rowman and Littlefield, Lanham, MD.

Pollock, J. L. (1967). Criteria and our knowledge of the material world. *The Philosophical Review*, 76(1):28–60.

Pollock, J. L. (1970). The structure of epistemic justification. *American Philosophical Quarterly*, (4):62–78.

Pollock, J. L. (1974). *Knowledge and Justification*. Princeton University Press.

Pollock, J. L. (1987). Defeasible reasoning. *Cognitive Science*, 11(4):481–518.

Pollock, J. L. (1991). A theory of defeasible reasoning. *International Journal of Intelligent Systems*, 6(1):33–54.

Priest, G. (1979). The logic of paradox. *Journal of Philosophical Logic*, 8(1):219–241.

Priest, G. (1994). The structure of the paradoxes of self-reference. *Mind*, 103(409):25–34.

Priest, G. (2006). *In contradiction*. Oxford University Press.

Prior, A. N. (1955). Curry's paradox and 3-valued logic. *Australasian Journal of Philosophy*, 33(3):177–182.

Pryor, J. (2013). Problems for credulism. In Tucker, C., editor, *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism*.

Recanati, F. (2004). *Literal meaning*. Cambridge University Press.

Reiter, R. (1980). A logic for default reasoning. *Artificial intelligence*, 13(1-2):81–132.

Restall, G. (1993). How to be really contraction free. *Studia Logica*, 52(3):381–391.

Restall, G. (1994). *On logics without contraction*. PhD thesis, University of Queensland.

Restall, G. (2002). *An introduction to substructural logics*. Routledge.

Restall, G. (2005). Multiple conclusions. In *Logic, methodology and philosophy of science: Proceedings of the twelfth international congress*, pages 189–205. Kings College Publications.

Ripley, D. (2013). Paradoxes and failures of cut. *Australasian Journal of Philosophy*, 91(1):139–164.

Ripley, D. (2015). Contraction and closure. *Thought: A Journal of Philosophy*, 4(2):131–138.

Ripley, D. (2017). Bilateralism, coherence, warrant. In *Act-Based Conceptions of Propositional Content: Contemporary and Historical Perspectives*, pages 307–324. Oxford University Press.

Ripley, D. (2018). On the 'transitivity' of consequence relations. *Journal of Logic and Computation*, 28(2):433–450.

Rosenblatt, L. (2019). On structural contraction and why it fails. *Synthese*, pages 1–26.

Ross, W. D. (1930). *The Right and the Good.Some Problems in Ethics*. Clarendon Press.

Rysiew, P. (2021). Epistemic Contextualism. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Spring 2021 edition.

Savage, L. (1954). *The foundations of statistics*. John Wiley and Sons.

Scanlon, T. M. (1998). *What We Owe to Each Other*. Harvard University Press, Cambridge, MA.

Scharp, K. (2013). *Replacing truth*. OUP Oxford.

Searle, J. R. (1969). *Speech acts: An essay in the philosophy of language*. Cambridge university press.

Searle, J. R. (1983). *Intentionality: An Essay in the Philosophy of Mind*. Cambridge University Press, Cambridge.

Shapiro, L. (2011). Deflating logical consequence. *The Philosophical Quarterly*, 61(243):320–342.

Shapiro, L. (2013). Validity curry strengthened. *Thought: A Journal of Philosophy*, 2(2):100–107.

Shapiro, L. (2015). Naive structure, contraction and paradox. *Topoi*, 34(1):75–87.

Shapiro, L. and Beall, J. (2018). Curry's Paradox. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, summer 2018 edition.

Simmons, K. (1993). *Universality and the liar: An essay on truth and the diagonal argument*. Cambridge University Press.

Slaney, J. (2010). A logic for vagueness. *The Australasian Journal of Logic*, 8.

Sosa, E. (1964). The analysis of'knowledge that p'. *Analysis*, 25(1):1–8.

Sosa, E. (1969). Propositional knowledge. *Philosophical Studies*, 20(3):33–43.

Sosa, E. (1970). Two conceptions of knowledge. *The Journal of Philosophy*, 67(3):59–66.

Sperber, D. and Wilson, D. (1995). *Relevance: communication and cognition*. 2 edition.

Stalnaker, R. C. (1968). A theory of conditionals. In *Ifs*, pages 41–55. Springer.

Standefer, S. (2016). Contraction and revision. *Australasian Journal of Logic*, 13(3).

Strasser, C. and Antonelli, G. A. (2018). Non-monotonic logic. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University.

Sturgeon, S. (2014). Pollock on defeasible reasons. *Philosophical Studies*, 169(1):105–118.

Sudduth, M. (2017). Defeaters in epistemology. In *Internet Encyclopedia of Philosophy*.

Swain, M. (1972a). An alternative analysis of knowing. *Synthese*, pages 423–442.

Swain, M. (1972b). Knowledge, causality, and justification. *The Journal of Philosophy*, 69(11):291–300.

Swain, M. (1974). Epistemic defeasibility. *American Philosophical Quarterly*, 11(1):15–25.

Tarski, A. (1936). Der wahrheitsbegriff in den formalisierten sprachen. *Studia Logica*, 1:261–405.

Titelbaum, M. G. (2019). Reason without reasons for. *Oxford Studies in Metaethics*, 14.

Travis, C. (2000). *Unshadowed thought: Representation in thought and language*. Harvard University Press.

Travis, C. (2008). *Occasion-sensitivity: Selected essays*. Oxford University Press.

Weber, Z. (2014). Naive validity. *The Philosophical Quarterly*, 64(254):99–114.

Whittle, B. (2004). Dialetheism, logical consequence and hierarchy. *Analysis*, 64(4):318–326.

Williamson, T. (1996). Knowing and asserting. *Philosophical Review*, 105(4):489.

Zardini, E. (2011). Truth without contra(di)ction. *Rev. Symb. Log.*, 4(4):498–535.

Zardini, E. (2019). Instability and contraction. *Journal of Philosophical Logic*, 48(1):155–188.