

FROM LOGICAL EXPRESSIVISM TO EXPRESSIVIST LOGIC: SKETCH OF A PROGRAM AND SOME IMPLEMENTATIONS¹

Robert Brandom
University of Pittsburgh

I. Introduction

Traditionally, two principal issues in the philosophy of logic are the demarcation question (what distinguishes specifically *logical* vocabulary?) and the correctness question (what is the *right* logic?). One of the binding-agents tying together semantic and logical inferentialism is a distinctive philosophy of logic: logical expressivism. This is the view that the expressive role that distinguishes logical vocabulary is to make explicit the inferential relations that articulate the semantic contents of the concepts expressed by the use of ordinary, nonlogical vocabulary. If one offers this logically expressivist, semantically inferentialist answer to the demarcation question, the correctness question lapses.

It is replaced by a concrete task. For each bit of vocabulary to count as logical in the expressivist sense, one must say what feature of reasoning, to begin with, with *nonlogical* concepts, it expresses. Instead of asking what the *right* conditional is, we ask what dimension of normative assessment of implications various conditionals make explicit. For instance, the poor, benighted, and unloved, classical two-valued conditional makes explicit the sense of “good inference” in which it is a good thing if an inference does *not* have true premises and a false conclusion. (At least we can acknowledge that implications that do *not* have at least this property are bad.) Intuitionistic conditionals in the broadest sense let us assert that there is a

¹ The proof-theoretic logical systems I report on in this paper were developed as the result of many years of work in our logic working group at the University of Pittsburgh, brought to fruition by Ulf Hlobil and Dan Kaplan. We will present them, along with many more, including some by Shuhei Shimamura, in the co-authored book we are writing, *Logics of Consequence: Tools for Expressing Structure*.

procedure for turning an argument for the premises of an inference into an argument for the conclusion. C.I. Lewis's hook of strict implication codifies the sense in which it is a good feature of an inference if it is *impossible* for its premises to be true and its conclusion not to be true. And so on. There can in principle be as many conditionals as there are dimensions along which we can endorse implications.

In spite of its irenic neutrality concerning the correctness question, one might hope that a new approach to the philosophy of logic such as logical expressivism would not only explain features of our *old* logics but ideally also lead to *new* developments in logic itself. I think this is in fact the case, and I want here to offer a sketch of how.

II. Prelogical Structure

I take it that the task of logic is to provide mathematical tools for articulating the *structure* of reasoning. Although for good and sufficient historical reasons, the original test-bench for such tools was the codification of specifically *mathematical* reasoning, the expressive target ought to be reasoning generally, including for instance and to begin with, its more institutionalized species, such as reasoning in the empirical sciences, in law-courts, and in medical diagnosis.

We can approach the target-notion of the structure of reasoning in two stages. The first stage distinguishes what I will call the “relational structure” that governs our reasoning practices. Lewis Carroll's fable “Achilles and the Tortoise” vividly teaches us to distinguish, in John Stuart Mill's terms, “premises from which to reason” (including those codifying implication relations) from “rules in accordance with which to reason,” demonstrating that we cannot forego the latter wholly in favor of the former. Gil Harman sharpens the point in his argument that there is no such thing as rules of deductive reasoning. If there were, presumably a paradigmatic one would be: If you believe p and you believe *if p then q* , then you should believe q . But that would be a terrible rule. You might have much better reasons against q than you have for either of the premises. In that case, you should give up one of them. He concludes that we should distinguish

relations of implication, from *activities of inferring*. The fact that p , *if p then q* , and *not- q* are incompatible, because p and *if p then q* stand in the implication relation to q , normatively *constrains* our reasoning activity, but does not by itself *determine* what it is correct or incorrect to do.

The normative center of reasoning is the practice of assessing reasons for and against conclusions. Reasons *for* conclusions are normatively governed by relations of *consequence* or *implication*. Reasons *against* conclusions are normatively governed by relations of *incompatibility*. These relations of implication and incompatibility, which constrain normative assessment of giving reasons for and against claims, amount to the first significant level of *structure* of the *practice* of giving reasons for and against claims.

These are, in the first instance, what Sellars called “*material*” relations of implication and incompatibility. That is, they do not depend on the presence of logical vocabulary or concepts, but only on the contents of non- or pre-logical concepts. According to semantic inferentialism, these are the relations that articulate the conceptual contents expressed by the *prelogical* vocabulary that plays an essential role in formulating the premises and conclusions of inferences.

Once we have distinguished these *relations* from the practice or activity of reasoning that they normatively govern, we can ask after the *algebraic* structure of such relations. In 1930s, Tarski and Gentzen, in the founding documents of the model-theoretic and proof-theoretic traditions in the semantics of logic, though differing in many ways in their approaches (as Jarda discusses in the second half of his book), completely agree about the algebraic structure of *logical* relations of consequence and incompatibility. Logical consequence satisfies Contexted Reflexivity (or Containment), Monotonicity, and Idempotence (Gentzen’s “Cut”, sometimes called “Cumulative Transitivity”). In Tarski’s terms: $X \subseteq \text{Cn}(X)$, $X \subseteq Y \Rightarrow \text{Cn}(X) \subseteq \text{Cn}(Y)$, and $\text{Cn}(\text{Cn}(X)) = \text{Cn}(X)$. Logical incompatibility satisfies what Peregrin calls “explosion”: the implication of *everything* by logically inconsistent sets. (Peregrin builds this principle in so deeply that he takes the functional expressive role of negation to be serving as an “explosion detector.”)

Perhaps these are, indeed, the right principles to require of specifically *logical* relations of consequence and incompatibility. But logical expressivists must ask a prior question: **What is the structure of *material* relations of consequence and incompatibility?** This is a question the tradition has not thought about at all. But the answer one gives to it substantially shapes the logical enterprise when it is construed as expressivism does.

We can think of statements of implication and incompatibility as expressing what is *included in* a premise-set and what is *excluded by* it. In a semantic inferentialist spirit, we can say that the elements of a premise-set are its *explicit* content, and its consequences are its *implicit* content—in the literal sense of what is *implied by* it. It is reasonable to suppose that what is *explicitly* contained in a premise-set is also part of its *implicit* content. It is accordingly plausible to require that material consequence relations, no less than logical ones, be reflexive in an extended sense: if the premises explicitly contain a sentence, they also implicitly contain it, regardless of what other auxiliary premises are available. (We sometimes call this condition "Containment", thinking of Tarski's algebraic closure principle that every premise-set is a subset of its consequence-set.)

Monotonicity, by contrast, is *not* a plausible constraint on *material* consequence relations. It requires that if an implication (or incompatibility) holds, then it holds no matter what additional auxiliary hypotheses are added to the premise-set. But outside of mathematics, almost all our actual reasoning is *defeasible*. This is true in everyday reasoning by auto mechanics and on computer help lines, in courts of law, and in medical diagnosis. (Indeed, the defeasibility of medical diagnoses forms the basis of the plots of every episode of "House" you have ever seen—besides all those you haven't.) It is true of subjunctive reasoning generally. If were to I strike this dry, well-made match, it would light. But *not* if it is in a very strong magnetic field. Unless, additionally, it were in a Faraday cage, in which case it would light. But *not* if the room were evacuated of oxygen. And so on.

The idea of "laws of nature" reflects an approach to subjunctive reasoning deformed by a historically conditioned, Procrustean ideology whose shortcomings show up in the need for idealizations (criticized by Cartwright in *How the Laws of Physics Lie*) and for "physics

avoidance” (diagnosed by Wilson in *Wandering Significance* on the basis of the need to invoke supposedly “higher-level” physical theories in *applying* more “fundamental” ones). Defeasibility of inference, hence nonmonotonicity of implication relations, is a structural feature not just of probative or permissive reasoning, but also of dispositive, committive reasoning. *Ceteris paribus* clauses do not magically turn nonmonotonic implications into monotonic ones. (The proper term for a Latin phrase whose recitation can do *that* is “magic spell.”) The expressive function characteristic of *ceteris paribus* clauses is rather explicitly to *mark* and *acknowledge* the defeasibility, hence nonmonotonicity, of an implication codified in a conditional, not to cure it by *fiat*.

The logical expressivist (including already—as I’ve argued elsewhere—Frege in the *Begriffsschrift*, at the dawn of modern logic) thinks of logical vocabulary as introduced to let one *say* in the logically extended object-language what material relations of implication and incompatibility articulate the conceptual contents of logically atomic expressions (and, as a bonus, to express the relations of implication and incompatibility that articulate the contents of the newly introduced logical expressions as well). There is no good reason to restrict the expressive ambitions with which we introduce logical vocabulary to making explicit the rare material relations of implication and incompatibility that are monotonic. Comfort with such impoverished ambition is a historical artifact of the contingent origins of modern logic in logicist and formalist programs aimed at codifying specifically *mathematical* reasoning. It is to be explained by appeal to historical causes, not good philosophical reasons.

Of course, since our tools were originally designed with this task in mind, as we have inherited them they are best suited for the expression of monotonic rational relations. But we should not emulate the drunk who looks for his lost keys under the lamp-post rather than where he actually dropped them, just because the light is better there. We should look to shine light where we need it most.

Notice that reasons *against* a claim are as defeasible in principle as reasons *for* a claim. Material incompatibility relations are no more monotonic in general than material implication relations. Claims that are incompatible in the presence of one set of auxiliary hypotheses can in

some cases be reconciled by suitable additions of collateral premises. Cases with this shape are not hard to find in the history of science.

What about Cut, the principle of cumulative *transitivity*? It is expressed in Tarski's algebraic metalanguage for consequence relations by the requirement that the consequences of the consequences of a premise-set are just the consequences of that premise-set, and by Gentzen as the principle that adding to the explicit premises of a premise-set something that is already part of its implicit content does not add to what is implied by that premise-set.

Thought of this way, Cut is the dual of what is usually thought of as the weakest acceptable structural principle that must be required if full monotonicity is not.² “Cautious monotonicity” is the structural requirement that adding to the explicit content of a premise-set sentences that are already part of its implicit content not defeat any implications of that premise-set. (Even though there might be *some* additional premises that *would* infirm the implication, sentences that are *already implied* by the premise-set are not among them.)

We can think generally about the structural consequences of the process of *explicitation* of content, in the sense of making what is *implicitly* contained in (or excluded by) a premise-set, in the sense of being implied by it, *explicit* as part of the explicit premises. Cut says that explicitation never *adds* implicit content. Cautious monotonicity says that explicitation never *subtracts* implicit content. Together they require that ***explicitation is inconsequential***. Moving a sentence from the right-hand side of the implication-turnstile to the left-hand side does not change the consequences of the premise-set. It has no effect whatever on the implicit content, on what is implied. (Explicitation can also involve making explicit what is implicitly *excluded* by a premise-set.)

² On holding onto both Cut and Cautious Monotonicity, see Gabbay, D. M., 1985, “Theoretical foundations for nonmonotonic reasoning in expert systems”, in K. Apt (ed.), *Logics and Models of Concurrent Systems*, Berlin and New York: Springer Verlag, pp. 439–459. Gabbay agrees with the criteria of adequacy laid down by the influential KLM approach of Kraus, Lehman, and Magidor: Kraus, Sarit, Lehmann, Daniel, & Magidor, Menachem, 1990. Nonmonotonic Reasoning, Preferential Models and Cumulative Logics. *Artificial Intelligence*, 44: 167–207.

Explicitation in this sense is not at all a *psychological* matter. And it is not even yet a strictly *logical* notion. For even *before* logical vocabulary has been introduced, we can make sense of explicitation in terms of the structure of *material* consequence relations. Noting the effects on implicit content of adding as an explicit premise sentences that were already implied is already a process available for investigation at the semantic level of the *prelogic*.

It might well be sensible to require the inconsequentiality of explicitation as a structural constraint on *logical* consequence relations. But just as for the logical expressivist there is no good reason to restrict the rational relations of implication and incompatibility we seek to express with logical vocabulary to monotonic ones, there is no good reason to restrict our expressive ambitions to consequence relations for which explicitation is inconsequential. On the contrary, there is every reason to want to use the expressive tools of logical vocabulary to investigate cases where explicitation *does* make a difference to what is implied.

One such case of general interest is where the explicit contents of a premise-set are the records in a *database*, whose implicit contents consist of whatever consequences can be extracted from those records by applying an *inference engine* to them. (The fact that the “sentences” in the database whose material consequences are extracted by the inference engine are construed to begin with as *logically* atomic does not preclude the records having the “internal” structure of the arbitrarily complex datatypes manipulated by any object-oriented programming language.) It is by no means obvious that one is obliged to treat the results of applying the inference-engine as having exactly the same epistemic status as actual entries in the database. A related case is where the elements of the premise-sets consist of experimental *data*, perhaps measurements, or observations, whose implicit content consists of the consequences that can be extracted from them by applying a *theory*. In such a case, explicitation is far from inconsequential. On the contrary, when the CERN supercollider produces observational measurements that confirm what hitherto had been purely theoretical predictions extracted from previous data, the transformation of rational status from *mere* prediction *implicit* in prior data to actual empirical observation is an event of the first significance—no less important than the observation of something incompatible with the predictions extracted by theory from prior data. This is the very nature of empirical *confirmation* of theories. And it often happens that

confirming *some* conclusions extracted by theory from the data infirms *other* conclusions that one otherwise would have drawn.

Imposing Cut and Cautious Monotonicity as global structural constraints on material consequence relations amounts to equating the epistemic status of premises and conclusions of good implications. But in many cases, we want to acknowledge a distinction, assigning a lesser status to the products of risky, defeasible inference. In an ideal case, perhaps this distinction shrinks to nothing. But we also want to be able to reason in situations where it is important to keep track of the difference in status between what we take ourselves to know and the shakier products of our theoretical reasoning from those premises. We shouldn't build into our global structural conditions on admissible material relations of implication and incompatibility assumptions that preclude us from introducing logical vocabulary to let us talk about those rational relations, so important for confirmation in empirical science.

The methodological advice not unduly to limit the structure of rational relations to which the expressive ambitions of our logics extend applies particularly forcefully to the case of incompatibility relations. The structural constraint the classical tradition for which Gentzen and Tarski speak imposes on incompatibility relations is *explosion*: the requirement that from incompatible premises anything and everything follows. This structural constraint corresponds to nothing whatsoever in ordinary reasoning practices, not even as institutionally codified in legal or scientific argumentative practices. It is a pure artifact of classical logical machinery, the opportune but misleading translation of the two-valued conditional into a constraint on implication and incompatibility that reflects no corresponding feature of the practices that apparatus—according to the logical expressivist—has the job of helping us to talk about. It is for that reason a perennial embarrassment to teachers of introductory logic, who are forced on this topic to adopt the low invocations of authority, pressure tactics, and rhetorical devices otherwise associated with commercial hucksters, con men, televangelists, and all the other sophists from whom since Plato we have hoped to distinguish those who are sensitive to the normative force of the better reason, whose best practices, we have since Aristotle hoped to codify with the help of logical vocabulary and its rules. In the real world, we are often obliged to reason from sets of premises that are explicitly or implicitly incompatible. [An extreme case is the legal practice of

“pleading in the alternative.” My defense is first, that I never borrowed the lawnmower, second, that it was broken when you lent it to me, and third that it was in perfect condition when I returned it. You have to show that *none* of these things is true. In pleading this way I am not confessing to having assassinated Kennedy. Examples from high scientific theory are not far to seek.] We should not impose structural conditions in our *prelogic* that preclude us from *logically* expressing material relations of incompatibility that characterize our actual reasoning. Explosion is not a plausible structural constraint on material relations of incompatibility, and our logic should not require us to assume that it is.

III. The Expressive Role of Basic Logical Vocabulary.

The basic claim of logical expressivism in the philosophy of logic is that the expressive role characteristic of *logical* vocabulary is to make explicit, in the object-language, relations of implication and incompatibility, including the material, prelogical ones that, according to semantic inferentialism, articulate the conceptual contents expressed by nonlogical vocabulary, paradigmatically ordinary empirical descriptive vocabulary. The paradigms of logical vocabulary are the *conditional*, which codifies relations of implication that normatively structure giving reasons *for* claims, and *negation*, which codifies relations of incompatibility that normatively structure giving reasons *against* claims.

To say that a premise-set implies a conclusion, we can write in the metalanguage: “ $\Gamma \vdash A$ ”. To say that a premise-set is incompatible with a conclusion, we can write in the metalanguage “ $\Gamma, A \vdash \perp$ ”.

To perform its defining expressive task of codifying implication relations in the object language, conditionals need to satisfy the

Ramsey Condition: $\Gamma \vdash A \rightarrow B$ iff $\Gamma, A \vdash B$.

That is, a premise-set implies a conditional just in case the result of adding the antecedent to that premise-set implies the consequent. A conditional that satisfies this equivalence can be called a “Ramsey-test conditional,” since Frank Ramsey first proposed thinking of conditionals this way.

To perform its expressive task of codifying incompatibility relations in the object language, negation needs to satisfy the

Minimal Negation Condition: $\Gamma \sim \neg A$ iff $\Gamma, A \sim \perp$.

That is, a premise-set implies not-A just in case A is incompatible with that premise-set. (It follows that $\neg A$ is the minimal incompatible of A, in the sense of being implied by everything that is incompatible with A.)

We should aspire to expressive logics built onto material incompatibility relations that are nonmonotonic as well as material implication relations that are nonmonotonic. That means that just as an implication $\Gamma \sim A$ can be defeated by adding premises to Γ , so can an incompatibility. Sometimes, $\Gamma, A \sim \perp$ can also be defeated, the incompatibility “cured”, by adding some additional auxiliary hypotheses to Γ . And while, given the role negation plays in codifying incompatibilities, an incompatible set, $\Gamma \cup \{A\}$ that is, one such that $\Gamma, A \sim \perp$) will imply the negations of all the premises that are its explicit members, it need not therefore imply *everything*. In substructural expressive logics built on Gentzen’s multisuccedent system, the condition that emerges naturally is not *ex falso quodlibet*, the classical principle of explosion, but what Ulf Hlobil calls “*ex fixo falso quodlibet*” (EFF). This is the principle that if Γ is not only materially incoherent (in the sense of explicitly containing incompatible premises) but *persistently* so, that is *incurably, indefeasibly* incoherent, in that *all* of its supersets are also incoherent, *then* it implies everything. In a monotonic setting, this is equivalent to the usual explosion principle. In nonmonotonic settings, the two conditions come apart. One conclusion that might be drawn from expressive logics is that what mattered all along was always *ex fixo falso quodlibet*—classical logic just didn’t have the expressive resources to distinguish this from explosion of all incoherent sets.

The basic idea of expressivist logic is to start with a language consisting of nonlogical (logically atomic) sentences, structured by relations of material implication and incompatibility. In the most general case, we think of those relations as satisfying the structural principles *only* of extended reflexivity—not monotonicity, not cautious monotonicity, and not even transitivity in the form of Cut. We then want to introduce logical vocabulary on top of such a language. This

means extending the language to include arbitrarily logically complex sentences formed from the logically atomic sentences by repeatedly applying conditionals and negations, and then extending the underlying material consequence and incompatibility relations to that logically extended language in such a way that the Ramsey condition and the Minimal Negation Condition both hold. (In fact, we'll throw in conjunction and disjunction as well.)

A basic constraint on such a construction is set out by a simple argument due to Ulf Hlobil.³ He realized that in the context of Contexted Reflexivity and a Ramsey conditional, Cut entails Monotonicity. For if we start with some arbitrary implication $\Gamma|\sim A$, we can derive $\Gamma, B|\sim A$ for arbitrary B —that is, we can show that arbitrary additions to the premise-set, arbitrary weakenings of the implication, preserves those implications. And that is just monotonicity. For we can argue:

$\Gamma \sim A$	Assumption
$\Gamma, A, B \sim A$	Contexted Reflexivity
$\Gamma, A \sim B \rightarrow A$	Ramsey Condition Right-to-Left
$\Gamma \sim B \rightarrow A$	Cut, Cutting A using Assumption
$\Gamma, B \sim A$	Ramsey Condition Left-to-Right.

Since we want to explore adding Ramsey conditionals to codify material implication relations that are reflexive but do not satisfy Cut—so that prelogical explicitation is not treated as always inconsequential—we will sacrifice Cut in the logical extension.

It is a minimal condition of logical vocabulary playing its defining expressive role that introducing it extend the underlying material consequence and incompatibility relations *conservatively*. (Belnap motivates this constraint independently, based on considerations raised by Prior's toxic “tonk” connective. The logical expressivist has independent reasons to insist on conservativeness: only vocabulary that conservatively extends the material relations of consequence and incompatibility on which it is based can count as *expressing* such relations explicitly.) So there should be no implications or incompatibilities involving only *old* (nonlogical) vocabulary that hold or fail to hold in the structure logically extended to include

³ Hlobil, U. (2016), “A Nonmonotonic Sequent Calculus for Inferentialist Expressivists.” In Pavel Arazim and Michal Dančák (eds.) *The Logica Yearbook 2015*, pp. 87-105, College Publications: London.

new, logical vocabulary, that do not hold or fail to hold already in the material base structure. Since that material base structure is in general nonmonotonic and intransitive, satisfying only contexted reflexivity, so must be the global relations of consequence and incompatibility that result from extending them by adding logical vocabulary.

IV. Basic Expressivist Logics

We now know how to do that in the context of Gentzen-style substructural proof theory. I will be summarizing technical work by recent Pitt Ph.D. Ulf Hlobil, now at Concordia University (on single-succedent systems) and current Pitt Ph.D. student Dan Kaplan (on multi-succedent systems).

We produce substructural logics codifying consequence and incompatibility relations that are not globally monotonic or transitive by modifying Gentzen's systems in three sequential stages. Gentzen's derivations all begin with what he called "initial sequents," in effect, axioms, (which will be the leaves of all *logical* proof trees) that are instances of immediate or simple reflexivity. That is, they are all of the form $A|\sim A$. We impose instead a structural rule that adds all sequents that are instances of *contexted* reflexivity—that is, (in the multisuccedent case) all sequents of the form $\Gamma, A|\sim A, \Theta$. Making this change does not really change Gentzen's system LK of classical logic at all. For he can derive the contexted version from immediate Reflexivity by applying Monotonicity, that is Weakening (his "Thinning"). So, as others have remarked, Gentzen does not need the stronger principle of unrestricted monotonicity in order to get the full system LK of classical logic. He can make do just with the *very* restricted monotonicity principle of Contexted Reflexivity, which allows arbitrary weakening *only* of sequents that are instances of reflexivity, that is, which have some sentence that already appears on both sides of the sequent one is weakening. Since all Gentzen's initial sequents are instances of immediate reflexivity, being able to weaken them turns out to be equivalent to being able to weaken *all* logically derivable sequents. (The weakenings can be "permuted up" the proof trees past

applications of connective rules in very much the same way Gentzen appeals to in proving his Cut-Elimination Hauptsatz.) Substituting the stronger version of Reflexivity for Gentzen's version accordingly allows dropping the structural requirement of Monotonicity. Contexted Reflexivity arises most naturally in Tarski's algebraic-topological way of thinking about consequence relations, as the principle that each premise-set is contained in its consequence set: $\Gamma \subseteq \text{con}(\Gamma)$.

We also do not impose Cut as a global structural constraint. But Gentzen's Cut-Elimination theorem will still be provable for all proof-trees whose leaves are instances of (now, contexted) Reflexivity. So the purely logical part of the system will still satisfy Cut.

The next step in modifying Gentzen's systems is to add axioms in the form of initial sequents relating logically atomic sentences that codify the initial base of *material* implications (and incompatibilities). Whenever some premise-set of atomic sentences Γ_0 implies an atomic sentence A , we add $\Gamma_0 | \sim A$ to the initial sequents that are eligible to serve as leaves of proof-trees, initiating derivations. (We require that this set of sequents, too, satisfies Contexted Reflexivity. We will be able to show that the connective rules preserve this property.) This is exactly the way Gentzen envisaged substantive axioms being added to his logical systems so that those systems could be used to codify substantive theories—for instance, when he considers the consistency of arithmetic. The crucial difference is that he required that these sequents, like those governing logically complex formulae, satisfy the structural conditions of Monotonicity and Cut—and we do not. We will introduce logical vocabulary to extend material consequence and incompatibility relations that do *not* satisfy Monotonicity, and that are *not* idempotent.

The third stage in modifying Gentzen's systems is accordingly to extend the pre-logical language to include arbitrarily logically complex sentences formed from that pre-logical vocabulary by the introduction of logical connectives. Gentzen's connective rules show how antecedent consequence and incompatibility relations governing the logically atomic base language can be systematically extended so as to govern the sentences of the logically extended language. Gentzen's own rules can be used to do this, with only minor tweaks. Like Ketonen's version of Gentzen's rules, ours are *reversible*. They are unlike the Gentzen-Ketonen rules in

that we mix additive and multiplicative rules. They are all equivalent to Gentzen’s own rules in the presence of a global structural rule of Monotonicity. But in nonmonotonic settings, they come apart. So, for instance, Gentzen’s left rule for conjunction allows us to move from $\Gamma, A | \sim C$ to $\Gamma, A \& B | \sim C$. That builds in monotonicity on the left. We can’t have that, since in the material base, it can happen that adding B as a further premise defeats the implication of C by Γ and A. We allow instead only the move from $\Gamma, A, B | \sim C$ to $\Gamma, A \& B | \sim C$. (A similar shift is needed in his right rule for disjunction: where he allows derivation of $\Gamma | \sim A \vee B, \Theta$ from $\Gamma | \sim A, \Theta$, building in monotonicity on the right, we allow instead only the move from $\Gamma | \sim A, B, \Theta$ to $\Gamma | \sim A \vee B, \Theta$.)

I said above that from a logical expressivist point of view, for the conditional to do its defining job of codifying implication relations in the object language, it needs to satisfy the Ramsey condition. In Gentzen’s setting, this amounts to the two principles:

$$\text{CP: } \frac{\Gamma, A | \sim B}{\Gamma | \sim A \rightarrow B} \quad \text{and} \quad \text{CCP: } \frac{\Gamma | \sim A \rightarrow B}{\Gamma, A | \sim B}$$

The first is Gentzen’s right-rule for the conditional. The second rule is not one of his. And it cannot be. For it is a simplifying rule. The only simplifying rule he has is Cut, and it is of the essence of his program to show that he can do without that rule: that every derivation that appeals to that single simplifying rule can be replaced by a derivation that does not appeal to it. Ketonen-style invertibility of connective rules, which makes root-first proof searches possible, though, requires not only Conditional Proof but the simplifying rule Converse Conditional Proof. And it is possible to show that this rule, too, like Cut is “admissible” in Gentzen’s sense: every derivation that uses it can be replaced by a derivation that does not.

It can be shown that our versions of Gentzen’s connective rules produce a *conservative extension* of any *nonmonotonic* material base consequence relation (including nonmonotonic incompatibility relations incorporated in such consequence relations) that satisfies the structural condition of Contexted Reflexivity. That is, in the absence of explicitly imposing a structural rule of Monotonicity (Weakening or Thinning) and Cut, the connective rules do not force global monotonicity. So the resulting, logically extended consequence relation is nonmonotonic. And the nonmonotonicity extends to logically complex formulae, for instance, as we have seen, in that from the fact that $\Gamma, A | \sim C$ it does not follow that $\Gamma, A \& B | \sim C$, so that from $\Gamma | \sim A \rightarrow C$ it does

not follow that $\Gamma|\sim(A\&B)\rightarrow C$. The logical language that results permits the explicit codification using ordinary logical vocabulary of arbitrary *nonmonotonic*, insensitive material consequence relations in which prelogical explicitation is *not* inconsequential.

And yet, the system is supraclassical. All the theorems of Gentzen’s system LK of classical logic can be derived in this system. For if we restrict ourselves to derivations all of whose leaves are instances of Contexted Reflexivity, that is, are of the form $\Gamma, A|\sim A, \Theta$, the result is just the theorems of classical logic. It is only if we help ourselves to initial sequents that are *not* of that form, the axioms that codify *material* relations of consequence and incompatibility, that we derive nonclassical results. **Gentzen never needed to require monotonicity, his “Thinning,” as a *global* structural rule. He could just have used initial sequents that correspond to Contexted Reflexivity instead of immediate reflexivity.** That gives him all the weakening behavior he needs. Further, if we look only at sequents that are derivable *no matter what material base relation we extend*, sequents such as $\Gamma, A, A\rightarrow B|\sim B$, hence $\Gamma|\sim(A\&(A\rightarrow B))\rightarrow B$, we find that the “logic” of our system in this sense, too, is just classical logic. Perhaps not surprisingly, if, following Gentzen, we use essentially the same connective rules but restrict ourselves to *single* succedent sequents, the result is a globally nonmonotonic, intransitive *supraintuitionist* logic.⁴

I have been talking about the logical extension of nonmonotonic material *consequence* relations and not about the logical extension of nonmonotonic material *incompatibility* relations. But the latter are equally well-behaved. The multi-succedent connective rules for negation are just Gentzen’s. But it is *not* the case that any materially incoherent premise-set implies every sentence. Such premise-sets imply both the sentences they explicitly contain and the negations of all those sentences. But they do not imply everything else. If a premise-set explicitly contains both A and $\neg A$ for some sentence A , *then* it implies everything. But that is because *persistently* or monotonically incoherent premise-sets explode—that is, sets that are not only incoherent themselves, but such that *every superset* of them is incoherent. This is what Ulf Hlobil calls “*ex fixo falso quodlibet*.” No specific stipulation to this effect needs to be made. It arises naturally

⁴ We do have to add some special rules, to make up for some of the things that happen on the right in the cleaner multisuccedent system.

out of the connective rules in the multisuccedent setting. If monotonicity held globally, *ex falso quodlibet* and *ex fixo falso quodlibet*” would be equivalent. Outside of derivations all of whose leaves are instances of contexted reflexivity, in our systems, they are not.

So in a clear sense, the *logic* is monotonic and transitive—indeed, classical or intuitionistic, depending (as with Gentzen) on whether we look at multi-succedent or single-succedent formulations—but the logically extended consequence and incompatibility relations *in general*, are not.⁵ The logic of nonmonotonic consequence relations is itself monotonic. Yet it can *express*, in the logically extended object language, the nonmonotonic relations of implication and incompatibility that structure both the material, prelogical base language, and the logically compound sentences formed from them, as they behave in derivations that substantially depend on the material base relations.

Substructural expressivist logics suitable for making explicit nonmonotonic, nontransitive material consequence and incompatibility relations are accordingly not far to seek. They can easily be built by applying to nonlogical axioms codifying those material relations of implication and incompatibility variants of Gentzen’s connective definitions that are equivalent to his under his stronger structural assumptions. It turns out that the relations of implication and incompatibility that hold in virtue of their logical form alone are still monotonic and transitive, even though the full consequence and implication relations codified by the logical connectives is not. So if you want Cut and Weakening, you can still have them—for purely *logical* consequence. Remember that from the point of view of logical expressivism, the point of introducing logical vocabulary is not what you can *prove* with it (what implications and incompatibilities hold in virtue of their logical form alone) but what you can *say* with it. Expressivist logics let us *say a lot* more than is said by their logical theorems.

⁵ When I talk about “the logic” here this can mean *either* the theorems derivable just from instances of Contexted Reflexivity (following Gentzen) or what is implied by every premise-set for every material base relation of implication and incompatibility that satisfies Contexted Reflexivity.

V. Codifying *Local* Regions of Structure: Monotonicity as a Modality

The master-idea of logical expressivism is that logical vocabulary and the concepts such vocabulary expresses are distinguished by playing a characteristic expressive role. They let us talk, in a logically extended object language, about the *material* relations of implication and incompatibility—what is a reason for and against what—that already articulate the conceptual contents of *nonlogical* vocabulary, as well as the *logical* relations of implication and incompatibility built on top of those material relations. Expressivist logics are motivated by the idea that we unduly restrict the expressive power of our logics if we assume that the global structural principles that have traditionally been taken to govern purely *logical* relations of consequence and inconsistency must be taken also to govern the underlying *material* consequence and incompatibility relations. So we don't presuppose Procrustean *global* structural requirements on the material relations of consequence and incompatibility we want to codify logically. Here is a further idea we have developed in what I am calling “expressivist logics.” Instead of imposing structural constraints globally, we relax those conditions and introduce vocabulary that will let us *say explicitly*, in the logically extended object language, *that* they hold *locally*, wherever in fact they still do.

Material consequence relations, I have claimed, are not in general monotonic. But they are not always and everywhere *nonmonotonic*, either. *Some* material implications are *persistent*, in that they continue to hold upon arbitrary additions to their premises. It follows from the fact that the regular Euclidean planar polygon has more than three sides that its angles add up to more than 180°, no matter what additional premises we throw in. The mistake of the tradition was not to think that *there are* material implications like this, but to think that *all* material implications *must* be like this. Logical expressivists want to introduce logical vocabulary that explicitly marks the difference between those implications and incompatibilities that are persistent under the addition of arbitrary auxiliary hypotheses or collateral commitments, and those that are not. Such vocabulary lets us draw explicit boundaries around the islands of monotonicity to be found surrounded by the sea of nonmonotonic material consequences and incompatibilities.

From a Gentzenian perspective, expressivist logics work out a different way of conceiving the relations between *structure* and *connective rules*. Connectives are introduced to express local structures. The paradigm is the conditional, which codifies the implication turnstile, by satisfying the Ramsey condition in the form of CP and CCP. Conjunction codifies the comma on the left of the turnstile, and disjunction codifies the comma on the right of the turnstile (in multi-succedent systems). (Note that in our nonmonotonic setting, this requires multiplicative rather than additive rules for conjunction on the left and disjunction on the right.⁶) Negation codifies incompatibility (in Gentzen’s multisuccedent systems elegantly captured in the *relation* between commas on the left and commas on the right). Our expressivist logics show how, in addition to the structures already captured by traditional connectives, further connectives can be introduced to mark local regions of the consequence relation where structure such as monotonicity and transitivity hold. I’ll try to give some idea of how this works by sketching what is for us the paradigm case: monotonicity.

The first idea is to extend the expressive power of our proof-theoretic metalanguage, so as to be able to distinguish persistent implications. In addition to the generally nonmonotonic snake turnstile “ \vdash ”, we can introduce a variant with an upward arrow, “ \vdash^{\uparrow} ” to mark persistent implications, that is, those that hold monotonically. To do this is to add *quantificational* expressive power to our proof-theoretic metalanguage. $\Gamma \vdash^{\uparrow} A$ says that not only does Γ imply A , but so does every superset of Γ : $\Gamma \vdash^{\uparrow} A$ iff $\forall X \subseteq L[\Gamma, X \vdash A]$.

All the connective rules can then be stipulated to have two forms: one for each turnstile. So we can write the right-rule (CP) for our Ramsey-test conditional showing the persistence arrow as optional, as:

$$\frac{\Gamma, A \vdash^{\uparrow} B}{\Gamma \vdash^{\uparrow} A \rightarrow B}.$$

If there is no upward arrow on the top turnstile, then there is none on the bottom either. But if there is a persistence-indicating upward arrow on the premise-sequent, then there is one also on the conclusion sequent. If Γ together with A *persistently* implies B —no matter what further premises we adjoin to them—then Γ *persistently* implies the conditional—no matter what further

⁶ $\frac{\Gamma, A, B \vdash \Theta}{\Gamma, A \& B \vdash \Theta}$ and $\frac{\Gamma \vdash A, B, \Theta}{\Gamma \vdash A \vee B, \Theta}$ rather than $\frac{\Gamma, A \vdash \Theta}{\Gamma, A \& B \vdash \Theta}$ $\frac{\Gamma, B \vdash \Theta}{\Gamma, A \& B \vdash \Theta}$ and $\frac{\Gamma \vdash A, \Theta}{\Gamma \vdash A \vee B, \Theta}$ $\frac{\Gamma \vdash B, \Theta}{\Gamma \vdash A \vee B, \Theta}$.

premises we adjoin to it. That follows from the original rule, together with the definition of persistence.

From this more structurally relaxed, nonmonotonic vantage point, traditional monotonic logic looks just the way it would if there were a notationally suppressed upward arrow on all of its turnstiles.

Incompatibility (and so logical inconsistency) also looks different in this setting. We now can distinguish materially incoherent premise-sets, where $\Gamma|\sim\perp$, from *persistently* incoherent premise-sets. These are premise-sets that are not only incoherent, but whose incoherence cannot be cured by the addition of further premises. And we can restrict explosion to those *persistently* incoherent sets. If $\Gamma|\sim\perp$, then for any $A\in\Gamma$, $\Gamma|\sim A$ and $\Gamma|\sim\neg A$. But it need *not* follow that for arbitrary B , $\Gamma|\sim B$. That follows only if $\Gamma|\sim^\uparrow\perp$. In the single-succedent case, we stipulate this: *ex falso quodlibet* but *ex **fixo** falso quodlibet*: ExFF. In the multi-succedent case, we do not need this stipulation. It falls out of the standard Gentzen treatment of negation. Here we want to say that what was always right about the idea that everything follows from a contradiction (and in our systems, if $A\in\Gamma$ and $\neg A\in\Gamma$, then Γ is *persistently* incoherent, and *does* imply everything) is that *persistently* incoherent premise-sets imply everything. It's just that in rigidly monotonic systems, *all* incoherence is treated as persistent, so in that expressively impoverished setting, ExF and ExFF are equivalent.

Once the dual-turnstile apparatus is in place in the metalanguage, we can introduce a modal operator in the object language to let us say there *that* an implication holds persistently. The basic idea is to introduce a monotonicity-box that says that $\Gamma|\sim\Box A$ iff $\Gamma|\sim^\uparrow A$, that is, if and only if $\forall X\subseteq L[\Gamma, X|\sim A]$. To say that Γ implies $\Box A$ is just to say that Γ *persistently* (that is, monotonically) implies A . The monotonicity box is clearly a strong modality, in that if Γ implies $\Box A$, then it implies A . And it is an S4 modality, in that if Γ implies $\Box A$, then it implies $\Box\Box A$.

From the point of view of a globally nonmonotonic implication relation in which local pockets of monotonicity are marked in the object language by implication of modally qualified

claims, the assumption of global monotonicity appears as what happens when one looks only at the monotonicity-*necessitations* of claims, ignoring anything not of the form $\Box A$.

In fact, we can do a lot better than what I have indicated so far. The expressivist idea is that the point of introducing logical vocabulary is to provide expressive resources that let one make explicit crucial local structural features of relations of implication and incompatibility—in the first instance, *material* relations of implication and incompatibility, and only as a sort of bonus the *logical* relations of implication and incompatibility that are built on top of them. From this point of view, what matters most is local persistence of some *material* implications. For it is these regions of local monotonicity in the material base relations of consequence and incompatibility that we want to be able to capture with a monotonicity-modal operator. Happily, it turns out that all we really need is an upward-arrow turnstile marking implications that can be weakened by the addition of arbitrary sets of logically *atomic* sentences. Our versions of Gentzen’s connective rules then guarantee that arbitrary weakenings by sets of logically complex formulae will be possible when and only when arbitrary weakening by sets of atoms is possible according to the underlying material base consequence relation.

In addition to implications whose persistence is underwritten by peculiarities of the underlying material consequence relation, there are implications of sentences prefaced by the monotonicity box that reflect logical relations induced by the connective definitions. Sentences like these—for instance, $\Box(A \rightarrow A)$ —do not depend on vagaries of the material implication relations.

A further innovation, pioneered by Ulf Hlobil for supra-intuitionistic single-succedent systems and by Dan Kaplan for supra-classical multiple-succedent systems, is the introduction of a much more powerful way of marking quantificational facts about sequents in the proof-theoretic metalanguage. (For simplicity, I’ll continue to use the single-succedent case.) Instead of introducing a simple upward arrow, as I have appealed to in my sketch, we introduce an upward arrow subscripted with a set of sets. $\Gamma \sim^{\uparrow X} A$ is defined as holding just in case for every set of sentences $X_i \in X$, $\Gamma, X_i \mid \sim A$. (In fact it suffices here, too, to restrict the values of X to sets of sets of logical *atoms* in the nonlogical material base language, but I put that complication aside here.) Then the set X

specifies a set of sets of sentences that one can weaken Γ with, while preserving the implication of A . That is, it marks a *range of subjunctive robustness* of the implication $\Gamma \sim A$. These are sets of sentences that can be added to Γ as collateral premises or auxiliary hypotheses without defeating the implication of A .

The underlying thought is that the most important information about a material implication is not whether or not it is monotonic—though that is something we indeed might want to know. It is rather under what circumstances it is robust and under what collateral circumstances it would be defeated. All implications are robust under *some* weakenings, and most are *not* robust under *all* weakenings. The space of material implications that articulates the contents of the nonlogical concepts those implications essentially depend upon has an intricate localized structure of subjunctive robustness and defeasibility. That is the structure we want our logical expressive tools to help us characterize. It is obscured by commitment to global structural monotonicity—however appropriate such a commitment might be for purely *logical* relations of implication and incompatibility.

Here, too, our variants of Gentzen’s connective definitions, as well as those for the monotonicity box, are so contrived as to ensure that it suffices to look at ranges of subjunctive robustness of implications that are restricted to the logical atoms governed by *material* relations of consequence and incompatibility. The more fine-grained control over ranges of subjunctive robustness offered by the explicitly quantified upward arrow apparatus is governed by a couple of structural principles. To indicate their flavor: one lets us combine sets of sets under which a particular implication is robust:

$$\frac{\Gamma \sim^{\uparrow X} A \quad \Gamma \sim^{\uparrow Y} A}{\Gamma \sim^{\uparrow X \cup Y} A} \quad \text{Union}$$

If the implication of A by Γ is robust under weakening by all the sets in X and it is robust under weakening by all the sets in Y , then it is robust under weakening by all the sets in $X \cup Y$. The very same connective rules stated with ordinary turnstiles go through as well with these quantified upward arrows with the same subjunctive-robustness subscript, and so propagate down proof trees.

The result of the addition of this apparatus is extensions of material consequence and incompatibility relations to a language including logically complex sentences, including those formed using the monotonicity modal box, that is well-defined and conservative of the material base relations. It follows that if the base relations are nonmonotonic and do not satisfy any version of Cut, then neither will the extended ones. The only structural principle we do impose on the base consequence relation, Contexted Reflexivity, is preserved. We do not impose the simplifying rule of Converse Conditional Proof (CCP)

$$\frac{\Gamma|\sim A\rightarrow B}{\Gamma, A|\sim B}$$

as a rule, but can prove it admissible, that is, as holding as a consequence of the connective rules for the conditional we do impose. The system is suprainuitionistic, in the single-succedent case, and supraclassical, in the multisuccedent case. If we restrict ourselves to elaborating material base consequence relations that consist entirely of instances of contexted reflexivity, that is of sequents of the form $\Gamma_0, p|\sim p$ for atomic sentences, then the logics over the extended languages are simply intuitionism and classical logic, respectively. These are obviously monotonic (so the monotonicity box is otiose), and Cut is, as usual, provably admissible.

VI. Conclusion

Construed narrowly, logical expressivism is a response to the demarcation question in the philosophy of logic. It suggests that we think of logical vocabulary and the concepts such vocabulary expresses as distinguished by playing a particular expressive role. The expressive task distinctive of logical vocabulary as such is to make explicit relations of consequence and incompatibility—to allow us to *say* what claims follows from other claims, and what claims rule out which others. Construed more broadly, logical expressivism invites us not to think about logic as having any autonomous subject matter—not *logical truth*, nor even *logical consequence*. Logic does not supply a canon of right reasoning, nor a standard of rationality. Rather, logic takes its place in the context of an already up-and-running rational enterprise of making claims and giving reasons for and against claims. Logic provides a distinctive *organ of self-consciousness* for such a rational practice. It provides expressive tools for talking and thinking, making claims, about the relations of implication and incompatibility that structure the giving of reasons for and against claims.

We should want those tools to be as broadly applicable as possible. The rational relations of material consequence that articulate the contents of nonlogical concepts are not in general monotonic. Good inferences can be inquired by adding new information. Indeed, offering finitely storable reasons typically *requires* that the implications we invoke be defeasible. Logic should not ignore this fact, nor even aim to rectify it. Logic should aim rather to codify even nonmonotonic, intransitive reasoning.

What I have here called “expressive logics” do that. The tweaks required to the proof-theoretic apparatus Gentzen bequeathed us for it to be capable of codifying nonmonotonic, even intransitive, reasoning are remarkably small. That fact tends to confirm the expressivist’s philosophical claims about what the *point* of logic has been all along. Expressive logics move beyond traditional logic not only in being built on antecedent relations of material consequence and incompatibility and in refusing to impose all but the most minimal global structural restrictions on those relations.⁷ They also introduce logical vocabulary that lets one express, in the logically extended language with its logically extended relations of consequence and incompatibility, *local* regions where structural conditions *do* hold. The paradigm is the introduction of a modal operator to mark the special class of monotonic implications, those that *can* be arbitrarily weakened with further collateral premises. (That turns out to include all those that hold in virtue of the meanings of the logical connectives alone). The benefits of treating monotonicity as a modality are many, and the costs are few. Treating logic as built on and explicating (elaborated from and explicative of) material relations of consequence and incompatibility offers another option besides substructural logics, when relaxing global structural constraints. One can introduce logical vocabulary to codify fine-grained *local* structures. These *monotonicity-modal* expressivist logics implement technically a central methodological principle of expressivist logics: don’t presuppose Procrustean *global* structural requirements on the material relations of consequence and incompatibility one seeks to codify logically. Instead, relax those global structures and introduce vocabulary that will let one *say explicitly*, in the logically extended object language, *that* they hold *locally*, wherever in fact they still do.

⁷ Of course not everyone—relevantists, for example—will agree that contexted reflexivity *is* minimal structure. So it should be admitted that this is a contentious description.

End