

AT-WILL RELATIONSHIPS: HOW AN OPTION TO WALK AWAY AFFECTS COOPERATION AND EFFICIENCY

ALISTAIR J. WILSON AND HONG WU

ABSTRACT. We theoretically and experimentally examine the effects from adding a simple, empirically relevant action to a repeated partnership, the option to walk away. Manipulating both the value of the outside option, and its relative distribution among the partners, we examine the behavior of human subjects in a repeated prisoners' dilemma. In particular, we examine the degree of cooperation and the form of punishments used. Our findings indicate that cooperation rates are broadly unaffected by the value of the common outside option, but that the selection of supporting punishments—in-relationship defections or walking-away—are dictated by individual rationality. In contrast to the symmetric results, when outside options for partners are asymmetric we find stark selection effects over cooperation, with the potential for very high and very low efficiency, dependent on the precise division rule.

1. INTRODUCTION

Dissolving a relationship is a familiar, easy-to-understand dynamic response, where the threat of a cessation of future interaction can be used to incentivize cooperation. It is clearly a force in many ongoing relationships of interest to economists: Workers quit firms that treat them badly; firms terminate workers that are unproductive. Couples petition for divorce if their marriages become unhappy. Consumers take away their business after a bad experience; firms refuse service to problem customer (expelling students, denying policy renewals). But in these examples participants also have access to and make use of in-relationship punishments: Workers strike or slow down; firms demote workers or withhold bonuses. Couples argue but eventually reconcile. Businesses win back customers with steep discounts after a complaint, while problem customers can pay more to atone (a donation to the school, temporarily higher premiums). In environments where mistakes or bad outcomes are inevitable despite the best efforts of all involved in-relationship punishments can be preferable, allowing for the possibility of forgiveness and a return to efficient cooperation where leaving the relationship does not.

Date: October, 2016.

Key words and phrases. Outside Options, Endogenous Termination, Dissolution clauses, Imperfect public monitoring, Dynamic games.

Wilson: University of Pittsburgh, Department of Economics, 230 Bouquet Street, Pittsburgh, PA; alistair@pitt.edu. *Wu:* Hong Kong Polytechnic University, School of Accounting and Finance, hhkpuwu@polyu.edu.hk. Our thanks to the following: Pedro Dal Bó, Mark Dean, John Duffy, Guillaume Fréchette, Johannes Hörner, Lise Vesterlund, Stephanie Wang, audiences at Brown, ESA, Stony Brook and SITE, as well as two anonymous referees and an associate editor who provided substantial feedback.

Whether or not dissolution is preferable to in-relationship punishment for the individual depends on many factors: *What is the in-relationship punishment's expected value and is recoordination on the efficient outcome possible?* That is, what is the value to remaining in the relationship for the punishment phase. *How severe are the dissolution frictions?* For example, legal costs for separation, losses on the sale of illiquid jointly held assets, reputation shocks from being fired/divorced, etc. *What is the expected value from starting afresh in a new relationship?* Here we must account for equilibrium selection in the population, and the possibly selected population available to rematch. *Are outside options fixed, and commonly available, or do they depend on the way the relationship ends?* The party breaking a contract might give up certain contractual rights, or alternatively, the party instigating might have access to a better outside option if they move more quickly.

There are many moving parts to this calculation, where the values both to remaining in the relationship and dissolving it are endogenous, leading to multitude of equilibrium possibilities. To get a handle on such a complex problem, our paper uses a laboratory study to examine selected outcomes as we *exogenously* vary the value of outside-options value. Using a simple in-relationship stage game—a prisoner's dilemma (PD) game with imperfect public monitoring—our experimental treatments alter the possibility/payoffs attained by walking away from a current partner.

In a static PD game, and in finitely repeated play, economic theory predicts an inefficient equilibrium outcome: both players defect, with this outcome Pareto dominated by joint cooperation. However, where interactions are repeated indefinitely, more-efficient equilibria are possible. Folk theorems indicate that with patient enough participants, all individually rational payoffs can be supported in equilibrium (Fudenberg and Maskin, 1986). Participants can bootstrap a cooperative outcome by conditioning their future behavior on the actions chosen this round. The ensuing indeterminacy in the theoretical prediction for ongoing relationships has been addressed by a growing body of experimental work analyzing equilibrium selection (see Dal Bó and Fréchette, forthcoming, for a meta analysis) with strong support for conditional cooperation when it is supportable as an equilibrium. Moreover, in experiments with imperfect monitoring—where actions in the stage game are unobserved with outcomes providing an imperfect signal—behavior is found to be lenient and forgiving (Fudenberg et al., 2012). Subjects require multiple bad outcomes to enter a punishment phase (lenience), while punishment durations are short followed by a return to cooperation (forgiveness). Because signals are imperfect, entering a punishment phase is unavoidable, eventually, even when all parties are fully cooperative. The selection of lenient and forgiving punishments therefore serves to increase efficiency, as inefficient punishment phases are entered less often and are shorter when they do occur. To this environment, our paper adds an intuitive, empirically relevant punishment device: ending the relationship. As terminating a relationship precludes forgiveness, if this punishment becomes focal, efficiency may be reduced.

Our first set of experiments examine *symmetric* environments, where each participant has access to the same in-relationship actions/payoffs and a common outside option. Here the fundamental tension is the interplay between incentive compatibility and efficiency for the inside- and outside-the-relationship punishments. A slight strengthening of the folk theorem (weakly undominated, individually rational payoffs) implies termination should never be used when its value is lower than the value of the in-relationship minmax (in a PD game,

the expected payoff from joint defection). Though its *presence* could conceivably alter equilibrium selection, theory predicts that we should *not* see relationships dissolve along the path. In contrast, when the dissolution payoff exceeds the in-relationship minmax the effect of a change in the outside option directly affects the individually rational payoffs. Dissolved relationships are now possible in equilibrium along the path of play.

Our experimental results are broadly in line with the above hypotheses. When outside options are far below the in-relationship minmax—in which region, we can think of termination as a form of costly punishment—we see very low rates of dissolution. When high, we see more substantial dissolution rates. However, deviating from the precise prediction, we observe large increases in termination use once outside options exceed the *minimal realization* subjects can guarantee within the relationship, rather than the expectation of this amount. That is, subjects begin to use termination as a costly punishment once its value is not stochastically dominated by joint defection. However, the use of termination when it is weakly dominated leads to expected losses overall, as the costs of the punishment are not compensated by an increased cooperation rate.

In terms of overall efficiency, our experimental results point to a non-monotonicity over the outside-option value. Through increased cooperation the *presence* of a low-payoff outside option does provide an efficiency boost in comparison to a treatment without a termination option, even though the actual use of termination is rare. However, once the outside option is not stochastically dominated by joint defection we observe much higher rates of termination. As dissolution is very inefficient when chosen, and we do not observe increased cooperation rates, we therefore find a sharp drop in efficiency. After this decline, subsequent increases to the outside option increase efficiency. Though rates of cooperation and dissolution remain flat, the resulting inefficiency from ending a relationship decreases in proportion to the outside-option value, so that overall efficiency increases.

Our first set of treatments address the *size* of the pie on dissolution common to all participants, and examine how changes to the individually rational action/payoff affect outcomes. Our second set of treatments instead examine the effects from changes to the *distribution* within the relationship on dissolution, and the precise rules that dictate who gets what. Here we are motivated by relational contracts specifying how assets/costs are divided on termination. Examples range from prenuptial agreements in marriages to LLP’s incorporation documents, consumer mortgages and residential leases through to contractual clauses that indicate severance and non-compete restrictions. While each party retains the right to unilaterally void the contract, they do so subject to the costs/benefits outlined for early termination. Our second treatment set examines three plausible asymmetric divisions on dissolution: (i) an environment where the party terminating receives the larger amount; (ii) an environment where the party who is being terminated gets the larger amount; and (iii) an environment where an independent arbitrator/judge diagnoses “blame” for the partnership’s dissolution, and assigns a larger payoff to the more-cooperative party.

Our findings in these asymmetric environments are stark. Rewarding the party ending the relationship leads to the vast majority of subjects choosing to terminate at the first opportunity, despite the expected payoff from joint termination being Pareto dominated by the in-relationship minmax. Relationships with this type of division rarely seem to get off square one in our experiments. In contrast, in what seems to be a more-common arrangement in leases and labor contracts, an asymmetric division rewarding the party being terminated

substantially reduces the use of termination, relative to comparable symmetric treatments. Subjects instead make much greater use of in-relationship punishments. However, once subjects begin defecting, they are less likely to return to cooperation than comparable symmetric treatments. In this parametrization we show that the asymmetry in payoffs on termination leads to much greater selection of inefficient in-relationship punishments.

Finally, our arbitrator treatment produces very high cooperation rates, leading to the most-efficient outcomes across our studied environments. Subjects select cooperative strategies at very high rates. Because subjects are highly cooperative, and lenient towards failures when they do occur, the observed dissolution rates are substantially lower than our high outside-option symmetric treatments. Though termination is very inefficient when chosen here, the asymmetric division rule produces a useful marriage of properties for the punishment: incentive compatibility/credibility for those with strong beliefs the other is free riding, but also strong punishment power for defectors being punished. The net effect is a substantial increase in overall efficiency even though the punishment itself is inefficient.

Our arbitrator treatment highlights how inefficiencies introduced into the process of dissolving a relationship can be useful for increasing ex ante efficiency. Taken together with the first-mover advantage treatment, the results illustrate how institutional features that affect the division at the end of a relationship can have large effects on both the likelihood of dissolution and the behavior within the relationship.

Below we briefly discuss the connection between our paper and the theoretical and experimental literature. Following this, Section 2 describes our core experimental design. Section 3 outlines our symmetric treatments and hypotheses, followed by Section 4 which reports the results from these treatments. Sections 5 and 6 present the design of our asymmetric experiments and outline the results, respectively. Finally, Section 7 summarizes the paper's main results and concludes.

1.1. Literature. The folk theorem for repeated games with discounting is articulated in Fudenberg and Maskin (1986), and shows that any feasible, individually rational payoff can be sustained in equilibrium when players are sufficiently patient.¹ Repeated games with *imperfect* public monitoring were initially studied with respect to cartel behavior, with firms receiving imperfect signals of the other cartel members' quantity decisions from the market via prices (Porter, 1983; Green and Porter, 1984). Further theory for imperfect monitoring was developed in Abreu et al. (1986, 1990) demonstrating the simple structures that can support optimal cooperation.

The closest theory paper to our environment is Radner et al. (1986) which examines partnerships game with imperfect monitoring.² However, given the focus on endogenous endpoints in our experiments (choosing to end the relationship) our environment is more technically a stochastic game (see Dutta, 1995, for a folk theorem result). That is, we require a state variable (here whether or not a player has terminated before the current round) that determines the particular stage game played, where the state transitions are endogenous, determined by players' actions. In many stochastic games the focus is on Markov perfect

¹For earlier work see references within Fudenberg and Maskin (1986), in particular Friedman (1971).

²Fudenberg et al. (1994) extends the folk theorems to infinitely repeated games with imperfect public monitoring, articulating a condition (pairwise identification) on the monitoring technology for the folk theorem to hold. The partnership game in Radner et al. (and the stage game in our experiments) fail this condition, and points on the feasible-payoff frontier are not attainable in equilibrium, even for $\delta \rightarrow 1$.

equilibrium (Maskin and Tirole, 2001) a subset of perfect equilibria where strategies are conditioned only on the current state variable. Given the simple binary nature of the state in our experiments, and the complete lack of agency in one state (inactive partnerships), our focus will instead be on a larger set of public-perfect equilibria (see Fudenberg and Tirole 1991).³

Theoretical treatments for dividing surplus on dissolution have tended to focus on maximizing the ex post efficiency through division of the partnership’s assets (Cramton et al., 1987; Preston McAfee, 1992). Instead, our paper’s focus is on ex ante efficiency, incorporating the effects of behavior within the relationship before dissolution occurs.⁴ Related to this focus on overall efficiency Comino et al. (2010) posit that firms might strategically omit more-explicit exit clauses in relational contracts to ensure costly litigation on early termination, and thereby increasing cooperation within the relationship to avoid this outcome (also see Li and Wolfstetter, 2010).

Early experiments on infinitely repeated games showed that cooperation is greater when it can be supported in equilibrium, but that subjects fail to make the most of the opportunity to cooperate (see Roth and Murnighan, 1978; Murnighan and Roth, 1983; Palfrey and Rosenthal, 1994). More recent experiments (Dal Bó, 2005; Aoyagi and Fréchette, 2009; Duffy and Ochs, 2009; Dal Bó and Fréchette, 2011) have provided more detail across different environments, and have contributed to the equilibrium-selection questions raised by folk theorems. A meta-study of these experiments (Dal Bó and Fréchette, forthcoming) finds that long-run cooperation is predictable through a basin of attraction index, and that three simple strategies well-represent the majority of play: *Always Defect*, the *Grim Trigger*, and *Tit-for-Tat*. One result for repeated play with imperfect public monitoring is that selected strategies are “lenient” in not retaliating after a single bad signal and to “forgiving” as they to return to cooperation after a short punishment phase (Fudenberg et al., 2012).⁵

Costly punishments in a repeated PD game are examined in Dreber et al. (2008), where the existence of a punishment option substantially increases cooperation rates, but not the average payoff of the group. This result is related to our treatments with symmetric but low outside options, where termination is similar to a costly punishment. However, for higher outside options termination takes on a distinctly different role, where its presence alters the equilibrium set.

A number of experimental papers have examined endogenous termination and rematching. For instance, Rand et al. (2011) examine the endogenous formation of partnerships, analyze cooperation levels in a structured network where players can change partners. Their paper similarly examines termination as a potential punishment/selection device, but the value of terminating is endogenously determined. They find that when subjects can update their network connections frequently, cooperation is maintained at a much higher level through endogenous selection. Independent of the present paper, Honhon and Hyndman (2015) and

³For an experimental examination of the Markov perfection restriction see Vespa and Wilson (2015) and citations therein.

⁴For an experimental examination of ex post dissolution see Brooks et al. (2010).

⁵See also Aoyagi and Fréchette (2009), which finds a positive connection between signal quality and cooperation, and Embrey et al. (2013) which examines renegotiation proofness.

TABLE 1. Stage-game payoffs

		Payoff, $r_i(a_i, Y)$		Pr $\{\mathcal{S} (a_i, a_j)\}$		Expectation, $u_i(a_i, a_j)$	
		Y:		a _j :		a _j :	
		S	F	C	D	C	D
a _i :	C	\$1.50	0	0.98	0.5	\$1.47	\$0.75
	D	\$2.50	\$1.00	0.5	0.1	\$1.75	\$1.15

Gaudeul et al. (2015) also examine the effects of termination in dilemmas, where our paper’s distinct contribution is in its systematic exogenous variation of the outside-option.⁶

2. EXPERIMENTAL DESIGN

2.1. Repeated Partnership Game. Our experimental game has two players engaged in a repeated joint-production task with imperfect public monitoring, similar to that in Radner et al. (1986). In every round t of their interaction, each partner i simultaneously chooses a private action $a_i^t \in \{C, D\}$. Given the resulting action profile $a^t = (a_1^t, a_2^t)$ a public signal $Y^t \in \{\mathcal{S}(\text{success}), \mathcal{F}(\text{failure})\}$ is realized and observed by both players, alongside their round payoff $r_i(a_i^t, Y^t)$. The probability of a *success* each round is a function of the selected actions, where the expected round payoff to partner i given the action profile (a_1, a_2) is

$$u_i(a_1, a_2) = \Pr \{\mathcal{S} | (a_1, a_2)\} \cdot r_i(a_i, \mathcal{S}) + (1 - \Pr \{\mathcal{S} | (a_1, a_2)\}) \cdot r_i(a_i, \mathcal{F}).$$

The partnership continues indefinitely under an exponential discount rate $\delta = \frac{4}{5}$. The discounted-average expected payoff for the partnership is defined as

$$(1) \quad W_i(\{a_1^t, a_2^t\}_{t=1}^{\infty}) = (1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} u_i(a_1^t, a_2^t).$$

Table 1 indicates the payoff realizations $r_i(a_i, Y)$ and the conditional success rates $\Pr \{\mathcal{S} | a\}$ chosen for our experiment. The chosen game can be thought of as two partners making choices over their individual effort levels into a jointly held venture (with C being high effort and D zero effort). The two partners equally split a \$5 firm revenue on success and \$2 revenue on failure. If either chooses to put in high effort they individually incur an additional \$1 cost. Though costly, higher effort produces an increased likelihood of a successful outcome: if both players expend effort there is a 98 percent chance the outcome is success; if both players put in low effort the probability of success is just 10 percent; if one exerts high effort and the other free rides the success probability is 50 percent.

The expected payoffs $u_i(a_i, a_j)$ are given in Table 1, which illustrates that the stage-game is a PD in expectation.⁷ The individually rational payoff for the stage game—an important parameter for many of our eventual hypotheses—is \$1.15, attained when both players defect,

⁶See also Hyndman and Honhon (2014), which examines termination options in a coordination game. They find that subjects show a slight preference for flexibility in ending the relationship but over-use the termination option, with too much sensitivity to short-term noisy outcomes.

⁷The game is therefore a prisoner’s dilemma where outcomes are simple lotteries. Under constant relative risk-aversion there are *no* risk parameters that would reverse the PD ordering on the stage-game payoffs. Under constant absolute risk aversion, the risk aversion coefficient would need to be between 1.5 and 2.4

the stage-game’s Nash equilibrium. However joint cooperating yields \$1.47 in expectation for each player, which is the most-efficient joint outcome for the partnership.

Our main experimental treatments modify the repeated game above by adding a third action to the game, which allows either partner to unilaterally dissolve the partnership. If either party chooses this third action, *T*(ermination), no further action choices are made by either partner in subsequent rounds, and both parties receive a fixed payoff for this and subsequent rounds. Our experiments focus on two types of termination payoffs: (i) *Symmetric*, where the players’ payoffs after termination are given by (π, π) ; and (ii) *Asymmetric*, where the two players receive the payoffs $(\bar{\pi}, \underline{\pi})$ after termination, with the assignment of the the higher payment $(\bar{\pi} > \underline{\pi})$ dictated by players’ actions.

2.2. Core Experimental Environment. The design is between subject: students once recruited are assigned to sessions with a fixed treatment: *the particular payoffs on termination*. Subjects participate in sessions of 12–16 subjects, at the start of which they are provided with instructions (a representative example is included in the appendix) which were read aloud.

As asking subjects to provide infinite choice sequences is infeasible, our experiment mirrors an infinitely-repeated game with payoff discounting through an exogenous, stochastically determined end point. That is instead of scaling down the payoffs received next round and onward by a factor $\delta = 4/5$, we scale down the probability of obtaining *any* additional amount in the next and subsequent rounds, retaining the same stakes.⁸ After every round of the game where the partnership is still accumulating payment, there is a $1/5$ probability that payment for the supergame will end, and $4/5$ that it continues. The agents expected *discounted-average* payoff from the supergame is therefore given by (1), as the probability of getting paid for round $t \in \{1, 2, \dots\}$ is given by δ^{t-1} . Subjects are paid the sum of their realized round payoffs $r_i(a_i^t, Y^t)$ or (after a dissolution has occurred) their round termination payoff π , up to round where the partnership’s payment exogenously stops.

One drawback to using a stochastic endpoint is that observed relationships can be very short, where short interactions offer limited power to assess the dynamic strategies. To increase the length of the observed partnerships in the experiment, we use a block design (cf. Fréchet and Yuksel, 2013, for a methodological discussion). Subjects are only informed on when/whether payment has ended after every block of five rounds. That is, at the end of every round the computer rolls a 100-sided die, common to all subjects in a session. The first round where the die-roll exceeds 80 is the last round for which we pay subjects. However subjects only observe outcomes from these rolls after rounds 5, 10, 15, etc. If all five rolls are less than or equal to 80 the game continues to another block of five, otherwise the partnership ends and payment is made up to the first round with an 80-plus die roll.

We will refer to each repeated partnership as a supergame. At the end of each supergame, subjects are randomly and anonymously rematched, and they then begin a new supergame. Sessions continue for at least an hour (excluding the time taken to read instructions), and the first supergame to exogenously end after the hour has elapsed marks the end of the

(for payoffs in cents) to induce a different ordinal game; consistent preferences at this level would require experimental subjects to take 50¢ for sure over an even gamble between zero and a million dollars.

⁸This method for implementing repeated games with no fixed horizon goes back to Roth and Murnighan (1978) and has been used extensively in experimental studies of dynamic behavior.

session. Subjects received payment for *two* randomly selected supergames, where they were also given a \$5 fixed payment.

Within a session each supergame has the same exogenously determined duration for all participants, whether the partnership is dissolved or not. Subjects cannot influence their time in the laboratory, nor can they increase their payoffs by playing more supergames. However, one potential concern we had was that subjects might not use the termination option, as they have no actions to take if they do. To mitigate this, our experimental design has each subject participating in two partnerships concurrently (this method is also used in Hauk and Nagel, 2001).⁹ To facilitate having two partners, our random-matching protocol randomly and anonymously forms subjects into a circle at the beginning of each new supergame. The subjects' two supergame partners are the session participants clockwise and counterclockwise from their position in the randomly formed circle. In this way, we minimize the ability a subject has to affect their clockwise partner through actions they take with their counterclockwise partner. In addition, all elements of the design are held constant across treatments except the main treatment variable, the availability of the termination option, and its payoff when chosen.

2.3. Experimental Session Summary. In total we conducted 20 experimental sessions at the Pittsburgh Experimental Economics Laboratory, with a total of 291 subjects, recruited from the University of Pittsburgh general subject pool. For each session we recruited 18 subjects and ran with at most 16 from those attending, with an average of 14.6 subjects per session. Subjects' earnings ranged from a minimum of \$5 to a maximum of \$77.75 (including their \$5 guaranteed show-up payment). Table 2 summarizes the particular experiments carried out and the number of sessions run for each, where we introduce details for our symmetric treatments in Section 3 and our asymmetric treatments in Section 5.

3. SYMMETRIC OUTSIDE OPTIONS

3.1. Design. Our first set of treatments consider an institution where the two partners receive the same fixed payoff π if the partnership is dissolved. Interpretations for this are a partnership with jointly held assets with a net value of $2 \cdot \pi$ (expressed as an annuity) on dissolution, where the partners equally split the proceeds, or that each partner anonymously re-enters some stationary matching market for a new partner with an expected outcome (net any rematching costs) of π .

Our symmetric treatments are broken into two sets of treatment. The first set provides the main comparative static with three treatments and three sessions for each treatment: (i) *No-T*, our baseline treatment where termination is *not available as an action*; (ii) *S-75*, a treatment where termination is available, but the payoff from choosing it is lower than the in-relationship minmax ($\pi = \$0.75$); (iii) *S-125*, a treatment where termination is available and the payoff from choosing it is higher than the in-relationship minmax ($\pi = \$1.25$).

The second treatment set provides single-session observations across a broader range of outside-option values. For each of these treatments we collect data from a single experimental session where the payoff π is in $\{\$0.85, \$0.95, \$1.05, \$1.15, \$1.35\}$, where we label the

⁹Referred to as *Partnership Red* and *Partnership Blue* within the experiment, both partnerships are affected by session-wide realizations for the exogenous end of payment. However, the probability of success/failure in each partnership depends only on the partnership-specific actions chosen, and are independent of all other outcomes in the experiment.

TABLE 2. Experiment Summary

Treatment			Sessions	Subjects	Supergames
Type	Term. Payoff	Label			
Symmetric (Depth)	None	<i>No T</i>	3	43	905
	\$0.75	<i>S-75</i>			
	\$1.25	<i>S-125</i>			
Symmetric (Breadth)	\$0.85	<i>S-85</i>	1	16	192
	\$0.95	<i>S-95</i>			
	\$1.05	<i>S-105</i>			
	\$1.15	<i>S-115</i>			
	\$1.35	<i>S-135</i>			
Asymmetric	\$1.25/\$0.75	<i>A-First</i>	2	27	465
	\$1.35/\$1.25	<i>A-Last</i>			
	\$1.25/\$0.75	<i>A-Judge</i>			
			20	291	5,799

treatments as *S-85* to *S-135*. Combining the data from the second set of treatments with the first, our design provides data on all multiples of \$0.10 between \$0.75 and \$1.35.

3.2. Theory and Hypotheses. The effects from the addition of a termination option to the partnership game is to change the sets of feasible and individually rational (IR) payoffs. The feasible expected payoff set is the convex hull of the expected stage-game payoffs in the *Active* state and the termination payoff vector (π, π) . To define individual rationality here we add an additional refinement to the standard definition. Because termination can be unilaterally imposed there are always weak Nash equilibria of the game where both parties terminate in round one, regardless of the value for π . However, when $\pi < u_i(D, D) = \$1.15$, choosing to dissolve the partnership is weakly dominated. This is true both for the stage-game and the supergame.¹⁰ To remove weakly-dominated outcomes, the focus of our hypotheses will be on individually rational outcomes of the game where both players use weakly undominated strategies.

For weakly undominated individual rationality (wIR) our hypothesis is dichotomous, dependent on whether the extra-relational payoff is below or above the intra-relational min-max, $\pi < u_i(D, D)$ and $\pi \geq u_i(D, D)$, respectively. In the first case leaving the partnership is weakly dominated, and the wIR payoff is the standard PD minmax achieved by joint-defection. In the second case, because termination can be unilaterally imposed, no discounted-average payoff less than π is individually rational. We therefore have the following hypotheses:

¹⁰When $\pi < u_i(D, D)$, any supergame strategy σ that chooses the action *T* after any history $h \in \mathcal{H}_T$ is weakly dominated by the strategy σ' that replicates σ everywhere else, but chooses *D* for all choices subsequent to *h*.

Rationality Hypothesis 1. *When $\pi < u_i(D, D)$, termination is never used, and the expected discounted-average payoff vector satisfies $(W_1, W_2) \geq (u_1(D, D), u_2(D, D))$*

Rationality Hypothesis 2. *When $\pi \geq u_i(D, D)$, joint-defection forever is not used, and the expected discounted-average payoff vector satisfies $(W_1, W_2) \geq (\pi, \pi)$.*

Hypothesis 1 is slightly more amenable to experimental test than Hypothesis 2, as it precludes termination as an observed action at any point on the path of play, whereas the defect action can still be chosen along the path when $\pi \geq u_i(D, D)$, but only in transition to subsequent cooperation or where the other player terminates that same round.¹¹

The two hypothesis are illustrated graphically in Figure 1. In each sub-figure the light-gray shape represents the set of feasible discounted-average expected payoffs, and the darker-gray region the set of wIR payoffs. Figure 1(A) illustrates the partnership game without a termination option; 1(B) the game with a termination payoff of $\pi < u_i(D, D)$ (in particular $\pi = \$0.75$); and 1(C) illustrates $\pi > u_i(D, D)$ (in particular $\pi = \$1.25$).

While the above two hypotheses are derived from a weak behavioral requirement (weakly undominated and individually rational outcomes) strengthening the requirement does not produce substantially more-precise hypotheses over payoffs. For example, consider instead hypotheses structured around a (fairly strong) *equilibrium* solution concept—weakly undominated, symmetric, stationary, public-perfect equilibria (wSSPPE, see the appendix for further details on specific equilibrium constructions). With the exception of *S-135*—which has joint-termination as the unique SPE outcome for $\delta = 4/5$ —all of our other treatments allow for wSSPPE with a discounted-average payoff of approximately \$1.45, which is only just below the maximal symmetric payoff of \$1.47. Measuring efficiency relative to the upper bound of \$1.47 and lower-bound of \$1.15 (joint defection), even a strong equilibrium restriction at best only rules out outcomes with greater than 95 percent efficiency.

The insubstantial further restriction possible for high-efficiency outcomes combines with the wIR payoff pair in all of our symmetric treatments being implementable as a wSSPPE. As such, our hypotheses based on wIR and feasibility provide similar upper and lower bounds on the efficiency of outcomes to those based on equilibrium. We therefore focus on the simpler and less restrictive wIR requirement.

While slackness in the Rationality Hypotheses reflect the inherent indeterminacy for theory in these games, our experiment’s greater aim is to explore the resulting behavior to understand selection. Motivated by the prior experimental literature, our results address the following selection questions: (i) Is the rate of cooperation affected by the outside option/individually rational action? (ii) Do the punishments used to support cooperation shift with the termination payoff? (iii) Are the punishments used more or less forgiving/lenient when termination becomes the individually rational action? We examine these questions with the following null hypotheses:

Selection Hypothesis 3 (Cooperation). *Cooperation rates are unaffected by the presence of a termination option, and do not vary as the value of the outside option increases/becomes individually rational.*

Selection Hypothesis 4 (Forgiveness/Lenience). *The frequency with which lenient/forgiving strategies are selected is unaffected by the termination payoff.*

¹¹For all values of $\pi > u_i(D, D)$ the asymmetric-strategy pair (*Always Terminate, Always Defect*) is weakly undominated SPE.

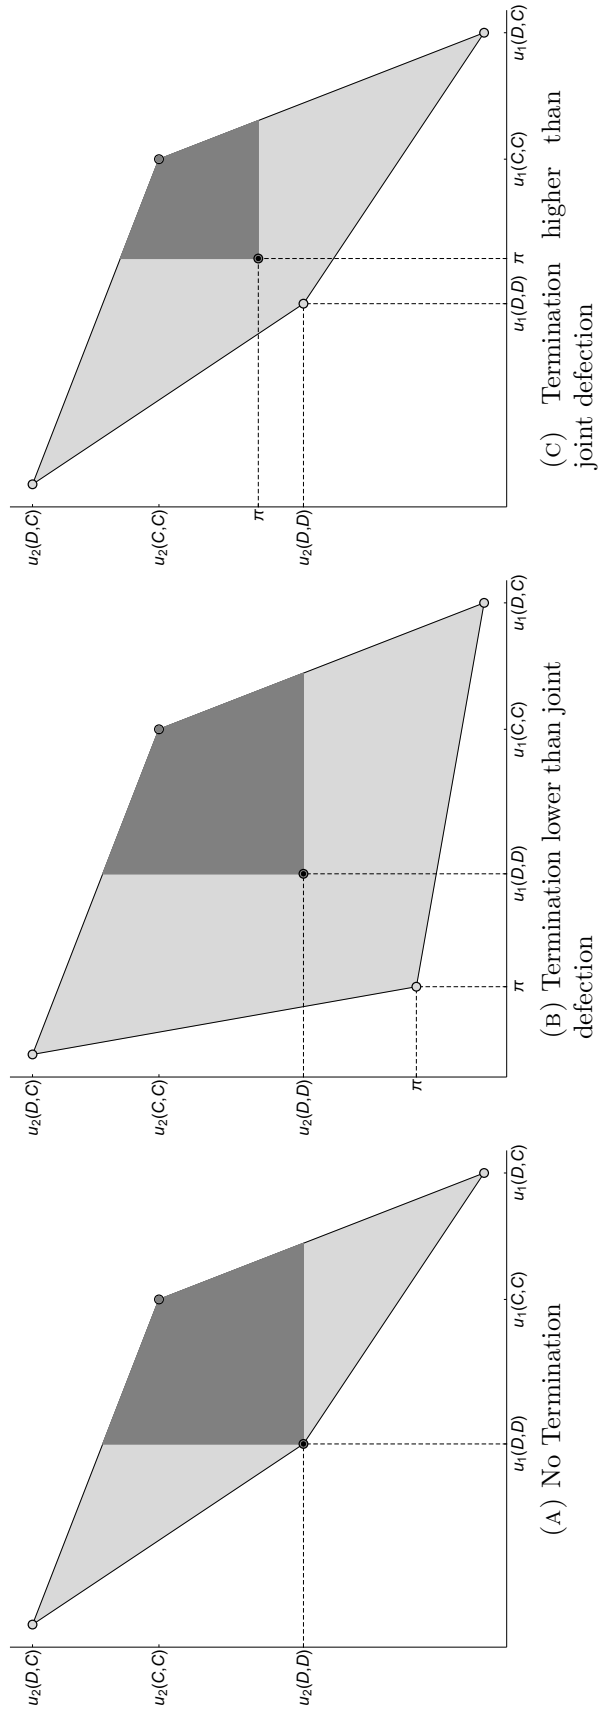


FIGURE 1. Sets of Feasible and Weakly Individually Rational Payoffs

Selection Hypothesis 5 (Punishments). *The punishments used to support cooperation stay constant as the outside option/individually-rational payoff changes.*

4. SYMMETRIC RESULTS

In terms of organization, our analysis of our symmetric results first examines aggregate choices and activity rates through sample averages, with the aim being to illustrate the broad patterns in the data. We then examine the evolution of behavior across sessions, analyzing how behavior changes as subjects gain experience within the session. We then illustrate the supergame dynamics, examining how subjects respond to the public signals choices within the supergame. The net effect of behavior on payoffs is analyzed in terms of payoff efficiency across treatments, where we comparing the results to our wIR hypotheses. Finally, we examine the selection of strategies with a more-structured estimation procedure, where we examine the effects of termination on our selection hypotheses.

Finding 1. *Aggregate data indicates that:*

- (A) *The presence of a dissolution option increases the rate of cooperation.*
- (B) *Activity rates for partnerships decrease with the outside option π .*
- (C) *As the termination payoff π increases there is no significant effect on cooperation in active relationships. However, unconditionally the cooperation rate does fall on the untruncated sample due to decreased activity.*

Evidence for the the above findings can be found in Table 3. Each cell in the table indicates a treatment-level average, where the table rows break-out each treatment and the data columns represent, from the second-most left column, moving to the right: (i) the cooperation rate in the first round of each supergame; (ii) the overall cooperation rate (assessed over rounds one to five, as all supergames have data for these rounds); (iii) the cooperation rate in the fifth round; (iv) the proportion of rounds where the relationship is active (rounds two to five only, as the first is active by construction); (v) the cooperation rate conditional on an active relationship (excluding data from inactive rounds where subjects have no agency); (vi) the fraction of uncooperative choices that use the termination option (so the residual amounts are defection choices); and finally columns (vii-viii) provide the number of subjects in each treatment (N_S) and the number of supergame observations (N_Γ), respectively.

Evidence for Finding 1(A), that the presence of a dissolution option increases cooperation, comes from a comparison of the treatment without a dissolution option (*No T*) to the symmetric-termination treatments as a group (*S-75-135*). Using session-level averages and a one-sided Mann-Whitney test we reject equivalence between the three *No-T* sessions and the eleven symmetric-termination treatments for first round cooperation ($p = 0.045$) and overall cooperation across the supergame’s first five rounds (also $p = 0.045$). That is, whether we look at how initially cooperative choices are or the average level across the partnership, the option to walk away from the partnership causes a significant increase in cooperation.

In terms of the size of this effect, the largest difference is in the initial cooperation rates. For the 1,810 *No-T* subject-supergame observations, 48.5 percent start with a subject choosing to cooperate, while this figure is 62.6 percent in the 6,632 subject-supergames with a symmetric

TABLE 3. Average Behavior

Treatment	Cooperation			Activity $2 \leq t \leq 5$	Active Choices		N_S	N_T
	$t = 1$	$1 \leq t \leq 5$	$t = 5$		C	$T \text{not } C$		
No T	0.485	0.417	0.361	1.000	0.417	0.000	43	905
S-75	0.628	0.569	0.512	0.970	0.583	0.020	43	974
S-85	0.630	0.543	0.487	0.980	0.552	0.019	16	192
S-95	0.828	0.708	0.625	0.941	0.743	0.070	16	384
S-105	0.632	0.467	0.371	0.749	0.584	0.127	16	368
S-115	0.512	0.443	0.373	0.792	0.532	0.100	16	288
S-125	0.602	0.472	0.368	0.758	0.586	0.153	42	774
S-135	0.533	0.408	0.311	0.741	0.515	0.139	16	336

Note: Treatments in bold have three sessions, all others single sessions. Number of subjects N_S , number of supergames N_T .

termination option. The presence of an option to walk away therefore leads to 14.1 percentage point increase in initial cooperation.

Looking at the overall cooperation rates in comparison to *No-T* it is helpful to break the symmetric treatments into two blocks, one block where the termination option is below a dollar (*S-75-95*, 5 sessions) and the other where it's above dollar (*S-105-135*, six sessions).¹² In both blocks the unconditional cooperation rate across the first five rounds of each supergame is significantly greater than the *No-T* treatment ($p = 0.072$ for *S-75-95* and $p = 0.083$ for *S-105-135*). However, while the size of the cooperation rate increase for *S-75-95* over *No T* is of a similar magnitude to that observed for first-round cooperation, for *S-105-135* the size of the gain is reduced below 5 percentage points. Moreover, for the *S-135* treatment the unconditional cooperation rate drops beneath the *No-T* average by a small amount, despite a higher rate of initial cooperation.

The reason for the smaller cooperation-rate gains in the termination treatments with high values of π can be seen by examining the *Activity* column in Table 2. The column indicates the proportion of supergame rounds 2-5 that are in the active state (excluding round one as it is active by default). The figure is 100 percent in the *No-T* treatment by construction, as termination is not an option here, but for treatments with termination the column indicates the extent to which either partner uses termination. For *S-75-95* the vast majority (94-98 percent) of supergames rounds are active, so subjects in low- π treatments use termination very infrequently. However, for the *S-105-135* treatments activity rates fall to 75-80 percent.

Where Finding 1(A) compares treatments with and without termination, Finding 1(B) stems from a comparison of the different values for π , establishing that termination use is responsive to the outside option. A least-squares regression of the 11 symmetric termination session activity rates on the termination payoff π indicates a significant negative effect

¹²We will subsequently use this division of the symmetric treatments about $\pi = \$1.00$ in further tests. We further motivate this point of division below.

($p = 0.002$) with an estimated drop in activity of 4.2 percent per \$0.10 increase in π . However, an alternative regression model where instead of π entering linearly we use the dummy $\mathbf{1}\{\pi > \$1\}$ as the RHS variable also finds a significant effect, but with better model fit.^{13,14} Regressing session activity on the dummy indicates an estimated activity drop of 19.7 percentage points once the termination payoff crosses the dollar threshold. Though the small number of session-level observations does not allow us to distinguish between the two models more precisely, the data indicates a qualitatively large decrease in activity once termination payoffs exceed a dollar.

Though the direction of Finding 1(B) is qualitatively in line with Hypothesis 1, the details are at odds with it. The hypothesis predicts full activity for all treatments where $\pi < \$1.15$. While for the S -75–95 treatments the inactivity rates are non-zero, they are close enough to the boundary that we can plausibly attribute termination use here to noise. However, for S -105 there is a large increase in the rate, and a quarter of rounds two to five are now inactive. One potential explanation for the discrepancy in S -105 is that the rationality hypotheses are formulated under risk-neutrality. While joint-defection yields a lottery of $\frac{9}{10} \cdot \$1 \oplus \frac{1}{10} \cdot \2.50 each round, choosing to terminate yields the certain value π . For subjects with enough risk aversion, termination might not be weakly dominated when $\pi = \$1.05$. A weaker form of Hypothesis 1 would require stochastic dominance of termination by the lottery $\frac{9}{10} \cdot \$1 \oplus \frac{1}{10} \cdot \2.50 , which gives precisely the regressor used in the second activity regression: no termination use whenever $\pi \leq \$1.00$.

Finally for our aggregate analysis, evidence for Finding 1(C) comes from a comparison of termination treatments with low outside-option values to those with a high value. Regressing session-level *active* cooperation rates on π there is no significant effect ($p = 0.713$). Similarly, a Mann-Whitney test comparing the S -75–95 and S -105–135 blocks finds no significant difference ($p = 0.204$). Where subjects have an active choice to make, the level of the outside option does not affect the likelihood with which a subject cooperates, all other things being equal.

While the value of termination in our experiments does not produce strong effects on the rate at which subjects choose to cooperate, an effect does emerge once we examine unconditional cooperation rates. That is, increasing the outside option does affect the rate at which partnerships cooperate in the long-run, but because of the increased number of dissolved partnerships when outside options are high. Looking at the unconditional cooperation rates in round five—our best experimental proxy for the long-run cooperation rate—the treatment averages in Table 3 reveal a clear pattern. While approximately a half or more of the partnerships in the S -75–95 treatments are still cooperative in round five, this figure is closer to 37 percent for S -105–135 ($p = 0.041$ for a Mann-Whitney test, $p = 0.037$ from a session-level regression). While the rate at which cooperation is selected each round is similar across our termination sessions, the cumulative effect from increased dissolution in S -105–135 is to reduce long-run cooperation. Instead of through the actions selected, reduced overall cooperation manifests through a form of sample censoring, as fewer partnerships survive to later rounds.

¹³This regression model with the dummy has an R^2 of 0.76, while the first model with a linear π term has an R^2 of 0.66. In terms of where we divide the data the division S -75–95 and S -105–135 is maximal with respect to the R^2 .

¹⁴Looking at session-level averages the activity rate is significantly greater in S -75–95 than S -105–135 ($p = 0.002$) using a Mann-Whitney one-sided test.




FIGURE 2. Initial Cooperation and Activity across Sessions

Note: Figures show the smoothed cooperation/activity (using a Laplace kernel) for partnerships within each session normalized to run from zero (the first supergame in the session) to one (the last supergame in each session).

Session Dynamics. Where the previous finding analyzed behavior in the aggregate, averaging across all session supergames, our next finding examines how outcomes evolve across experimental sessions.

Finding 2. *Across sessions we find:*

- (A) *Reduced cooperation in No T relative to the termination treatments emerges quickly.*
- (B) *Cooperation rates fall across symmetric termination sessions.*
- (C) *Activity rates decrease across sessions for $\pi > \$1.00$ treatments, and increase in the $\pi < \$1.00$ treatments.*

Evidence for the three statements in Finding 2 are illustrated within Figure 2. In Figure 2(A) we provide the cooperation rate in the first round of each supergame, where the horizontal axis represents time passing across the session. Similarly, Figure 2(B) presents information on the activity rate in round five of each supergame, with the horizontal axis again indicating progress across the session. Within each figure we show results for the *No-T* treatment (solid gray line), the *S-75–95* block of treatments (dash-dot black line) and the *S-105–135* block (dashed black line).¹⁵ In each sub-figure the lines represents a kernel-smoothed estimate for the average cooperation rate at this point in the session.

The first finding, apparent from Figure 2(A), is that reduced cooperation in the *No-T* treatment emerges early on within the session. The points that begin each figure indicate the cooperation rate at the beginning of each session, before any other interactions. While the symmetric termination treatments start at approximately 70 percent cooperation, the *No-T* treatment is significantly lower at 57 percent ($p = 0.019$, session level comparison). Initial cooperation rates in *No T* then decrease to approximately 46 percent mid-session before recovering to just over 50 percent cooperation at the end of the session.

¹⁵Results are qualitatively similar if we exclude the more-cooperative *S-95* session from the data, with the effect that the gap between the *S-75–85* and *S-105–135* treatments narrows.

The second result is the decrease in initial cooperation rates in the termination treatments as the session proceeds. The fall is largest in the S -105–135 treatments, falling to 54.4 percent initial cooperation in the last ten supergames of each session, where the S -75–95 treatments falls to 63.1 percent.¹⁶

Our final result across sessions comes from Figure 2(B), which shows the extent to which subjects make use of termination as the session evolves. Depicting the fraction of supergames in round five that are still active, the figure shows decreasing activity for the high-termination payoff block, and increasing activity rates (though less sharply) in the low-termination payoff block. Subjects learn not to use termination in the treatments where the payoff from doing so is stochastically dominated by staying in the partnership and defecting. For the treatments with $\pi > \$1.00$ we instead see decreasing activity across the session, as subjects learn to use termination more and more as they gain experience with the environment.¹⁷

Across sessions our results indicate that: (i) the *presence* of a dissolution option increases coordination on cooperation from the very first choice; however, (ii) the cooperation rates fall as subjects gain more experience; and (iii) subjects use termination at increased rates as the session proceeds when its value is not stochastically dominated by joint defection.

Supergame Dynamics. The previous section examined behavior *across* supergames, we now analyze behavior *within* supergames, assessing the extent to which subjects use a history-dependent response. We focus here on the aggregate-level behavior in response to the previous round, where a more-structured analysis of supergame *strategies* is provided at the end of this section.¹⁸

Finding 3. *Within supergame we find:*

- (A) *Across all treatments subjects' supergame response is history dependent, with continued cooperation less likely following a failure than a success.*
- (B) *Relative to No T, subjects in the termination treatments are more likely to continue cooperating, both in response to successes and failures in the last round.*
- (C) *Termination choices are most common in histories where the subject cooperated and observed a failure last round; however, termination rates following this history are only large in magnitude when termination is not stochastically dominated by defecting.*

Given an active partnership there are five possible histories from the previous round. The first is the empty history, \emptyset , which occurs in round one of each supergame, as nothing has happened within the partnership yet. The other four possible observable outcomes from the previous round are given by $\{C, D\} \times \{\mathcal{S}, \mathcal{F}\}$. That is, the subject knows their own action choice last round, as well as the public outcome of success or failure. In what follows we analyze subjects' choices in reaction to each of the five possible previous-round histories.

¹⁶A subject-random-effect probit on first round cooperation for each supergame finds a significant negative time trend as the session proceeds for both the S -75–95 and S -105–135 blocks ($p = 0.000$ for both), but no effect for the *No-T* treatment ($p = 0.458$).

¹⁷Time trends are significant in both the S -75–95 and S -105–135 blocks (both $p = 0.000$ from a random-effects probit). Looking at individual treatments the activity trends are negative and significant in S -105, S -125 and S -135; positive and significant in S -75 and S -85; and insignificant in S -95 and S -115.

¹⁸Given the above discussion on session dynamics, we restrict attention from this point forward to behavior in the last ten supergames in each session—the last five non-concurrent supergames, given the two partnerships. Subsequent results therefore reflect subjects that have gained experience with the task.




FIGURE 3. Active round choice, by last-round history (last ten supergames)

Figure 3 illustrates the history-dependence in behavior within active supergames. The three sub-figures illustrate the proportion of choices in active rounds for, respectively: (A) the *No-T* treatment; (B) the *S-75-95* treatments; and (C) the *S-105-135* treatments. In each sub-figure we show the first-round choices (following the history \emptyset) as the strip on the left, with the proportion of cooperation (shaded in white), defection (light gray) and termination (dark gray) indicated through the height of each bar. The larger box in each sub-figure is a unit-square representing all active subject-rounds beyond the first. We sub-divide this unit square along the horizontal direction according to the previous round's history and in the vertical direction by the chosen action.

The four possible histories from the previous round are therefore indicated by vertical strips in each sub-figure, where the width of each strip indicates the fraction of active rounds with this history. Within each vertical strip the proportion of choices for cooperate/terminate/defect are indicated by the heights of the white/dark-gray/light-gray bars, respectively. The height of each bar therefore indicates the sample analog to $\Pr \{a_t | h_{t-1}, \text{Active}\}$ while the width represents $\Pr \{h_{t-1} | \text{Active}\}$. By construction, the *area* of each entry therefore represents the joint probability $\Pr \{a_t, h_{t-1} | \text{Active}\}$ for all non-initial rounds.

Finding 3(A) is supported by the fact that subjects in all three treatment blocks cooperate much more following a round where they cooperated and observed a success than any of the other three options. This can be seen with the taller white bar after the history (C, S) than any of the other non-initial histories.¹⁹ Comparing the rate of cooperation following (C, S) to that observed after each of the three alternative histories we reject equivalence for *No T*, *S-75-95* and *S-105-135* (max of $p = 0.000$ across each of the three pairwise comparisons, in each treatment set).

Though all three treatment blocks have maximal cooperation in response to successful cooperation last round, there are certainly differences in the rate across the three figures.

¹⁹Statistical inference to complement the figure is carried out through a (subject level) random-effects probit estimate on the decision to cooperate in the sample of all active rounds except the first in the last ten supergames. The regressors are dummy variables for the four previous-round histories as the explanatory variables with no constant.

Finding 3(B) indicates that the rate of *continuing* cooperation is significantly greater in the treatments with termination than without.²⁰ This is true not only for continued cooperation after a success last round (with $p = 0.002$ for S -75–95, and $p = 0.042$ for S -105–135), but also after observing a failure (p -values of 0.004 and 0.001, respectively). This second comparison indicates that subjects seem to be more lenient in the termination treatments, with a greater likelihood of continued cooperation after observed failures.

Finally, Finding 3(C) focuses on the use of termination. In each of the Figure 3 illustrations termination choices are represented by the dark gray regions. In the two treatment blocks where termination is available, the decision to dissolve is most commonly made following a failed cooperation. Given (C, \mathcal{F}) last round, terminations represent 4.4 percent of choices in the S -75–95 treatments, and 27.2 percent of choices in the S -105–135 treatments. The only other history with non-negligible termination rates is the response to (D, \mathcal{F}) in the S -105–135 treatments, where 8.7 percent of the choices are to end the relationship.²¹

To examine the effect of π on the use of termination as a punishment, we use session-level averages from the last ten supergames. We calculate the termination rate following (C, \mathcal{F}) as our dependent variable and regress it on the session’s outside option π . Using the 11 symmetric termination session averages we find a significant positive effect from the outside option on the use of termination following failed cooperation. For every \$0.10 increase in the outside option the regression predicts a 5.6 percent increase in the use of termination (significantly different from zero with $p = 0.000$).

Payoffs and Efficiency. We now examine the efficiency of the selected outcomes, and test the hypotheses generated by individual rationality and weak dominance articulated in section 3. Our efficiency measure for each subject-supergame is the following transformation of the total discounted joint payoff:

$$\Upsilon \left(\{a_1^t, a_2^t\}_{t=1}^5 \right) = \frac{\frac{(1-\delta)}{(1-\delta^5)} \sum_{t=1}^5 \delta^{t-1} \frac{1}{2} (u_1(a_1^t, a_2^t) + u_2(a_1^t, a_2^t)) - u_1(D, D)}{u_1(C, C) - u_1(D, D)}.$$

That is, we take the discounted-average expected payoff to the two partners within the supergame, given the observed actions, and measure this relative to the in-relationship minmax payoff (\$1.15). This number is then normalized by the difference in expected payoff between mutual cooperation and mutual defection.

Finding 4. *With respect to payoffs and efficiency we find:*

- (A) *The weakly undominated, individually rational hypotheses (Rationality Hypotheses 1 and 2) are not rejected across symmetric treatments where cooperation is possible in equilibrium.*

²⁰For across-treatment comparisons we use a standard probit on the decision to cooperate or not in active rounds in the last ten supergames, with the RHS-variables reflecting all interactions of previous round history h and the treatment block (*No T*, S -75–95 and S -105–135). We do not include subject-level random effect as these would confound with treatment effects, so we instead cluster standard errors at the subject level.

²¹Within the S -105–135 treatments, the termination rate after (C, \mathcal{F}) is significantly larger than the rate after (D, \mathcal{F}) with $p = 0.000$. Both of these rates are significantly larger than the termination rate following rounds with a success ($p = 0.000$). For the S -75–95 treatments the termination after (C, \mathcal{F}) is significantly larger than after the other three histories (max p -value of 0.001).




FIGURE 4. Payoff Efficiency by treatment

Note: Shaded areas represent lower-bound from individual rationality hypotheses (achievable with *Always Defect/Always Terminate*) with the upper bound from the best SPE implementable with a symmetric, public-perfect three-state machine. Error bars indicate 95 percent confidence regions for average treatment payoff-efficiency calculated with a bootstrap of size 5,000 with supergame-level resampling.

- (B) *The presence of a termination option leads to more-efficient outcomes in the majority of our symmetric treatments.*
- (C) *Efficiency is not monotonic in the outside option, falling sharply as π increases above \$1.00. This fall comes about as termination is utilized more frequently as a punishment when $\pi > \$1.00$, but its use is Pareto dominated by in-relationship defection for $\pi \leq \$1.15$.*

We illustrate our efficiency measure Υ in Figure 4. For each of our symmetric treatments, the figure indicates the sample average for Υ across the last ten supergames in each session with a gray diamond, where the error bar indicates a 95-percent confidence region (calculated with a bootstrap). In addition to the overall discounted average, we also illustrate the average payoff efficiency for behavior in the first round only (a white circle) and fifth round only (a white triangle). The predictions made by Hypotheses 1 and 2 are overlaid on the figure as the gray shaded region.²²

Examining the efficiency levels relative to the wIR restriction on behavior, all of our treatments except *S-135* are significantly greater than the lower bound, both for the discounted-average and round one and round five in isolation). In contrast, for *S-135* we can reject the wIR hypothesis with 99 percent confidence.²³

²²The upper bound for the shaded regions are calculated from the best memory-one wSSPPE.

²³At $\delta = \frac{4}{5}$ the *unique* equilibrium payoff in *S-135* coincides with the IR payoff, as all equilibria involve at least one player terminating in the first round. The sample-average efficiency in the *S-135* session (54.6 percent) is significantly lower than the 62.5 percent prediction ($p = 0.002$). However, the round five efficiency in *S-135* (60.9 percent) is not significantly lower than the prediction ($p = 0.292$). Rounds one to four drive

Evidence for Finding 4(B) comes from regressing the discounted-average efficiency Υ within the last ten supergames on treatment dummies (with *No T* the omitted category). Four out of the seven symmetric termination treatments generate significantly higher efficiency than *No-T* (*S-75*, *S-95*, *S-125* and *S-135*, all below $p = 0.020$). We fail to reject equivalence for two treatments (*S-85* and *S-115*), while we find significantly lower efficiency in *S-105* than *No T* ($p = 0.005$).²⁴ While repeating this exercise for efficiency in the first-round only mirrors the results for the discounted average, regressing efficiency in the *fifth*-round on treatment dummies yields no significant differences from *No T* for any of the $\pi < \$1.00$ treatments.²⁵ This last result suggests that efficiency gains in the treatments where termination is stochastically dominated may be short-lived.

Our third finding relates to the specific pattern across the termination treatments. As illustrated by Figure 4, efficiency is not monotonic in π . The discounted-average efficiency increases as the outside option π approaches \$1.00, stemming from increased cooperation in the *S-95* treatment (in particular increased round-one cooperation). However, a sharp drop in efficiency occurs at $\pi = \$1.05$ as subjects begin to use the (at this point very inefficient) termination option at non-trivial rates. As π increases across the \$1.05–\$1.35 range efficiency increases too. However, unlike the changes to efficiency for $\pi < \$1.00$ that are driven by increased cooperation, the efficiency increases for $\pi > \$1.00$ are more-directly attributable to the changes in π . Cooperation and termination rates are somewhat constant for $\pi > \$1.00$, but the efficiency of a *dissolved* relationships is directly proportional to π . As such, efficiency increases across this region where termination is used.

Evidence for the fit for out non-monotonicity comes from a regression of the seven treatment-level averages for Υ in Figure 4 on the following three regressors: (i) a dummy for $\pi < \$1.00$; (ii) a dummy for $\pi > \$1.00$; and (iii) the interaction of the previous $\pi > \$1.00$ dummy with the difference $(\pi - \$1.05)$. While, this regression indicates a 46.8 percent average efficiency for the *S-75–95* treatments, this falls to 24.4 percent for the *S-105–135* treatment, but with a predicted marginal increase of 10.3 percentage points for each ten cent increase in the outside option over \$1.05. Though there are only seven treatment observations, the difference between the first and second regressor is significant ($p = 0.019$, from an *F*-test), as is the third regressor, the positive slope for efficiency from *S-105* to *S-135* ($p = 0.016$ from a *t*-test).²⁶

Our third finding, the non-monotonicity in π , is related to the results Dreber et al. (2008). When $\pi < u_i(D, D)$ termination effectively serves as a costly punishment, though here committing the partnership to the inefficient outcome for the remainder of the supergame.

the lower efficiency, reflecting a high degree of (C, D) and (D, D) choices (expected efficiencies of 31.3 percent and 0 percent, respectively).

²⁴Using all supergames instead of the last ten, *S-85* is also significantly greater than *No T* ($p = 0.000$) but both *S-105* and *S-115* are not significantly different from *No T* ($p = 0.291$ and $p = 0.251$, respectively).

²⁵The treatment effects for the fifth round efficiency—again relative to *No T* assessed over the last ten supergames—are similar to the discounted-average effects for treatments with $\pi > \$1.00$: *S-105* is significantly less efficient ($p = 0.000$) while *S-125* and *S-135* are significantly more efficient (both $p = 0.000$).

²⁶Using the same regression model to explain the first-round efficiencies (the white circles in Figure 4) we find no significant relationship (neither for the difference between the dummies, nor for the marginal effect on π above \$1.05), where this is explained by the very low rates of first-round termination. In contrast, for the fifth-round efficiency (the white triangles) the effects are both significant with the relevant *F/t* test ($p < 0.004$).

TABLE 4. Strategy Classifications (Last ten supergames)

	<i>No T</i>	<i>S-75-S-95</i>	<i>S-105-135</i>
Initial Cooperative	52.0%	67.5%	61.5%
Ongoing Cooperation	58.9%	65.8%	55.3%
Lenient Punishment	10.3%	35.9%	37.9%
Forgiving Punishment	28.4%	27.7%	26.3%
(+ <i>All-C</i>)	31.1%	46.2%	33.6%
Uses Termination	0.0%	5.1%	33.8%

Subjects in our experiments do not use this costly punishment when defection stochastically dominates termination. However, once the punishment is no longer stochastically dominated, subjects begin to use termination at more substantial rates, with the cost of substantially decreased payoffs. This mirrors the Dreber et al. result that though the presence of a costly punishment can be effective at increasing cooperation it is sub-optimal for subjects actually choosing to make use of it.

Strategy Estimation. Our final analysis for the symmetric treatments investigates the strategies adopted by subjects within each treatment. We use the method detailed in Dal Bó and Fréchette (2007), *Strategy Frequency Estimation Method* (SFEM, also used in Fudenberg et al., 2012 and Embrey et al., 2013).²⁷ To use SFEM we specify a set of 38 strategies, motivated by the previous experimental literature.²⁸ Given the strategy-set restriction, and an econometric error term (an independent probability of mistakes when implementing the strategy), we estimate the proportions of each strategy via maximum likelihood, using data from the last ten supergames in each. Appendix A outlines the method in more detail and reports full estimation results over the 38 strategies. We focus here on providing information on the broad families of strategy used, in particular whether the strategies are cooperative, forgiving, and/or lenient.

Finding 5. *With respect to strategy selection we find:*

- (A) *While the presence of a termination option does produce a small increase in the selection of strategies that start out cooperating, we do not detect economically or statistically large effects on the cooperation rate with respect to the value of the outside option (Selection Hypothesis 3).*
- (B) *The presence of a termination option leads to the selection of more-lenient strategies, but no strong effect on the selection of forgiving strategies (Selection Hypothesis 4).*
- (C) *The selection of strategies that use termination to punish failures increases significantly for the S-105-135 treatment block (Selection Hypothesis 5).*

²⁷In addition to this, a supplemental appendix provides a comparable reduced-form analysis of the data, looking at the one-round conditional response. This appendix effectively estimates an aggregate level memory-one public strategy (with probabilistic transitions).

²⁸We choose our strategy set to build upon those included in the previous literature. Taking the main strategy set included in Embrey et al. (2013), we add termination variants for strategies in those papers.

Evidence for the above Finding is given in Table 4, reporting the proportion of estimated strategies exhibiting: (i) initially cooperative behavior; (ii) the possibility for ongoing cooperation; (iii) lenience in response to failure; (iv) forgiving punishment phases; and (v) punishment phases with termination components.²⁹

In *No T*, where termination is not an option, the most-common estimated strategies are always defect (*All-D*), the grim-trigger (*Grim*) and a variant of tit-for-tat (called *Mono* in the literature) that starts out cooperating, but defects following failures and cooperates following successes in the previous round. Our baseline treatment therefore mirrors the meta-study strategy results in Dal Bó and Fréchette (forthcoming) that these three strategies account for the majority of play in infinite-horizon PD games.

For *No T* strategies that start out cooperating have a 52 percent weight in the estimations—which as a sanity check for the SFEM procedure closely matches the average first-round cooperation rate of 51 percent. Those strategies capable of sustaining cooperation with positive probability are represented by the *Ongoing Cooperation* category, which have an estimated incidence of 59 percent.

In comparison to *No T*, the estimated strategies for the *S-75–95* treatments do indicate an increase in the selection of strategies that start out cooperating. However the SFEM does not indicate this difference is significant when we compare the bootstrapped distributions ($p = 0.249$, the probability with which a *No T* estimate exceeds an *S-75–95* one for the sampling distributions). Similarly, despite an increase in the selection of strategies capable of ongoing cooperation to 66 percent the SFEM process does not indicate this difference is significant ($p = 0.259$) over the 59 percent rate estimated for *No T*.³⁰

Despite null results over the rates at which cooperative strategies are selected, we do find significant differences when we examine the rates at which subjects select lenient strategies. Comparing the low-outside-option treatments to *No T*, subjects do select more lenient strategies at greater rates when a termination action is available ($p = 0.048$). In contrast to lenience though, we do not find differences (neither in levels, nor inferentially) in the selection of forgiving strategies ($p = 0.447$).³¹

Comparing the *S-75–95* estimates those for *S-105–135*, the clearest difference is the more-frequent selection of terminating strategies in the treatments with high outside option values ($p = 0.085$). Our experimental results clearly point to the punishments used to support cooperation shifting with changes to the value of the outside option, reflecting the our individual rationality predictions. Grim-triggers, which are strategies selected at significant rates in our estimations for *No T* and *S-75–95* have no estimated incidence in the *S-105–135*

²⁹Possibility of ongoing cooperation is defined as a strategy that if cooperating in round m has a positive-probability of a sequence of public outcomes such that the strategy cooperates in all rounds $m + n$ for $n \in \mathbb{N}$. A lenient strategy is any strategy with a cooperative phase that transits to another cooperative state for both public outcomes (and so the strategy *Always Cooperate* is lenient). A forgiving strategy is a one with a defection phase which returns to a n ongoing cooperative phase state for some public outcome (we separately report the fraction of forgiving strategies when we classify *Always Cooperate* as forgiving). A strategy using termination is any machine which chooses the termination action with positive probability within a supergame.

³⁰We similarly fail to find significant effects when comparing the *S-105–135* treatments to either *No T* or *S-75–95* treatments.

³¹For the high-outside-option treatments the difference in selection of lenient strategies is only marginally significant ($p = 0.114$), despite the increased rate of selection in Table 4. The selection of forgiving strategies is not significantly lower than in *No T* ($p = 0.466$) or *S-75–95* ($p = 0.402$).

results. Instead of the grim triggers (which revert to defection as punishments) we find that subjects instead select strategies which cooperate initially but terminate after a number of failures (referred to as C- N -Strike, for the number of failures until termination).

While forgiving strategies (those capable of returning to cooperation after defecting) have a non-negligible incidence at 28 percent, lenient strategies (those that allow for more than one failure before shifting to defecting) are not common, with an estimated frequency of just 10.3 percent. and that the selected strategies are more lenient.

Taking the results from the SFEM exercise to our three selection hypotheses, while we do find evidence to reject the null in Selection Hypothesis 5 (the outside-option value *does* influence the selection of termination strategies), we do not find substantial evidence to reject the null for Selection Hypotheses 3 and 4. While the *presence* of a termination option does increase the selection of more-lenient strategies, we do not detect a clear pattern with respect to the *value* of the outside option. This result is repeated for the cooperation levels, where levels of initial cooperation are somewhat higher when we allow for a termination option,³² but the differences across termination treatments as we manipulate π are more muted.

Symmetric Treatment Summary. While there are certainly differences in cooperation rates across our treatments, the strongest effects seem to be driven by: (i) the presence of an option to terminate; (ii) a censoring effect where the use of termination precludes further cooperation in late-supergame rounds. Our interpretation of the first effect is that the presence of the termination option may help the subjects to think dynamically. However, beyond this, overall differences in cooperation rates are not substantial.

While cooperation rates do not differ as we vary the outside option π , the rate at which subjects make use of termination punishments does, and in a clearly predictable way. Where termination is stochastically dominated by remaining in the relationship we see very infrequent termination use, despite substantial miscoordination. Once termination is no longer stochastically dominated we see more substantial rates of inactivity as subjects start using termination as a punishment.

Where previous repeated PD experiments have shown that the selection of efficient cooperative strategies is related to payoff parameters such as the discount rate δ , and the relative risk/rewards of deviating from a cooperative outcome, our symmetric experiments show that the *form* of the individually rational action does not have a substantive effect on cooperation rates within an intermediate range. The main effect is instead on the form of punishment selected within uncooperative partnerships.

What seems to matter for selection of strategies is an intuitive and theoretically tractable feature, the individually rational payoff, as opposed to the specific features/labels/forms of the IR action (here, termination or defection). One interpretation of our symmetric results is that they provide evidence to support simple two-action PD games as useful reduced forms for more-complicated partnerships with differing forms of dynamic punishment devices.

Our symmetric treatments therefore tell a straightforward and theoretically simple story: While the selected punishments are dictated by the individually rational payoff, the effects on other aspects of the selection are more muted. Below we begin to paint a more-complex picture when outside options are asymmetric. Indeed, we show that when the division

³²See the aggregate Finding 1(B) and the conditional cooperation rates in Finding 3(B) for evidence of this outside the SFEM results.

of outside options are unequal and endogenous, cooperation and efficiency can be quite substantially different.

5. ASYMMETRIC TERMINATION PAYOFFS

Our asymmetric treatments examine the same repeated partnership game with a dissolution option. However the two partners now receive different payoffs on termination. One partner receives a high payoff amount $\bar{\pi}$, while the other gets a low payoff, $\underline{\pi}$.³³ In the three treatments we examine, which partner gets the higher payoff is endogenous. In each treatment we will make the focal payment for the participant choosing to terminate \$1.25.

In general the division rule on termination can depend on the entire history $\{a_1^t, a_2^t, y^t\}_{t=1}^T$ up to the point of termination. There is subsequently a vast constellation of asymmetric division rules. Our focus will be on three rules that in our view are both strategically interesting *and* empirically relevant.³⁴

Asymmetric First. Our first asymmetric treatment assigns the higher payoff to the party ending the relationship, and the lower payment to the party being terminated, where we label the treatment *A-First*. The treatment is motivated by partnerships where those that exit first are best prepared, or where last-movers are left stuck with large debts/costs. For examples, consider law and accountancy firms incorporated as partnerships that on hitting hard times saw large-scale partner defections to rival firms (with those leaving possibly taking many of their previous firm’s clients with them).³⁵ The *A-First* treatment mirrors this tension, where a strong belief that the other will leave the partnership makes leaving yourself a best response.

For this treatment we choose $\bar{\pi} = \$1.25$ and $\underline{\pi} = \$0.75$. The focal payoff for either partner *choosing* to terminate is \$1.25, as they will either get the high payoff or if both parties terminate that round they get $\hat{\pi} = \frac{1}{2}\bar{\pi} + \frac{1}{2}\underline{\pi} = \1.00 . Other than this change to the payoffs on termination, the game is identical to our symmetric treatments.

Unlike the symmetric environment, joint-termination in the very first round cannot be removed as an equilibrium outcome by appealing to weak dominance. This is true for any $\bar{\pi} > \underline{\pi}$, as termination is now the unique best response to the belief the other will terminate. Because $\hat{\pi} < \$1.15$ in this treatment, the expected payoff from joint termination is below that of joint defection. The treatment therefore introduces an equilibrium payoff pair $(\hat{\pi}, \hat{\pi})$ that is Pareto dominated by the in-relationship minmax of joint defection.

A second negative effect in *A-First* relative to a symmetric treatment with $\pi = \$1.00$ is that it becomes harder to support other equilibrium payoffs. Discounted-average payoffs in the interval $(\hat{\pi}, \bar{\pi})$ cannot be supported in equilibrium without a positive probability of termination, as otherwise either player could deviate and get the high termination payoff with certainty. So not only does *A-First* incorporate a Pareto-inferior equilibrium outcome,

³³All of our treatments with ties on the assignment-rule will break them with a fair coin, so the expected payoff on a tie is $\hat{\pi} = \frac{1}{2}\bar{\pi} + \frac{1}{2}\underline{\pi}$.

³⁴An interesting institution which we have not pursued is a fixed asymmetric division, where player one gets $\bar{\pi}$ with certainty on dissolution. Theoretically, this is not too different from the symmetric division discussion above, but can allow for ‘abusive’ equilibria where one player uses the threat of credible threat of termination to produce asymmetric outcomes such as (D, C) .

³⁵Arthur Andersen LLP, an accounting firm, saw large-scale partner defections after the Enron accounting scandal broke. Howrey LLP, a global law firm, similarly witnessed extensive partner defections to competing firms before its formal dissolution in 2011.

it also removes joint-defection forever as an equilibrium. Focusing on payoffs supportable with pure-strategies the predictions for *A-First* is:

Rationality Hypothesis 6. *In the A-First treatment the discounted-average payoff vector either satisfies $(W_1, W_2) > (\bar{\pi}, \bar{\pi})$ or involves joint termination in round 1 and $(W_1, W_2) = (\hat{\pi}, \hat{\pi})$.*

Asymmetric Last. Our second asymmetric treatment *A-Last* is the mirror of *A-First*, assigning the higher payoff to the party who has been terminated and the lower payoff to the partner choosing to terminate. Individual rationality is dictated by the point $\max\{\underline{\pi}, \min\{\bar{\pi}, u_i(D, D)\}\}$. That is, each player can guarantee themselves the higher of: (i) the lower termination payoff $\underline{\pi}$ by choosing to end the game themselves, or (ii) if they switch to defect forever, they can force their partner to choose between giving them the minimum of the joint-defection payoff and the high termination value $\bar{\pi}$.

This treatment was chosen to mirror contract clauses that force the party initiating early termination to incur additional costs. Many employment contracts allow for severance payments to executives (so called “golden parachutes”) if the firm voids the contract. If instead the relationship is voided by the employee the severance payment is not made (in the other direction, employees who voluntarily leave might have to adhere to non-compete agreements, reducing their outside options). One plausible effect of such clauses is that they can induce employees in faltering relationships to behave badly, hoping that the other party terminates them, rather than quitting themselves.

Unlike *A-First*, *A-Last* does not have *joint-termination* as an equilibrium outcome, as each partner does strictly better by defecting if they believe the other will terminate. However, similar to *A-First*, joint-defection forever is not supportable as an equilibrium outcome if $\underline{\pi} > \$1.15$. So for some parametrizations, cooperative outcomes cannot be supported by joint-termination, nor by joint-defection forever.

In order to focus on this tension our treatments use the parametrization $\underline{\pi} = \$1.25$ and $\bar{\pi} = \$1.35$.³⁶ Unilaterally dissolving the partnership is better than joint-defection forever, but each partner strictly prefers that the other party does the terminating. The individually rational payoff is \$1.25, where the wIR payoff set is the same as in the *S-125* treatment, which serves as a natural comparison treatment.³⁷

Rationality Hypothesis 7. *In the A-Last treatment the discounted-average payoff vector satisfies $(W_1, W_2) \geq (\underline{\pi}, \underline{\pi})$*

Asymmetric Judge. Our final asymmetric treatment is motivated by arbitration-hearings after a relationship ends. A judge/arbitrator (through some perfect, possibly costly, forensic process) obtains access to the complete history of play $\{a_1^k, a_2^k, Y^k\}_{k=1}^t$. After acquiring this information, she assigns the higher dissolution payoff to the party that cooperated most. That is, once any player chooses to end the relationship, the judge examines the

³⁶If $\underline{\pi} < U_i(D, D)$ then the wIR set in the *A-Last* treatment is identical to the symmetric termination games with $\pi_T < U_i(D, D)$ and termination is a strictly dominated choice. Our design therefore chooses the parametrization of $(\underline{\pi}, \bar{\pi}) = (\$1.25, \$1.35)$ for this treatment, given the more-compelling strategic tensions, rather than the values $(\$0.75, \$1.25)$ used in our other asymmetric treatments.

³⁷The equilibria of the *A-Last* game differ from *S-125*, which we discuss in more detail in the strategy appendix.

action sequence $\{a_1^k, a_2^k\}_{k=1}^t$ and assigns the higher termination payoff to player i and the lower payoff to player j if

$$\sum_{k=1}^t (\mathbf{1}\{a_i^k = C\} - \mathbf{1}\{a_j^k = C\}) > 0.$$

We label this treatment *A-Judge*, as it is intended to mirror an institution where a third party divides partnership assets taking into account the partner’s behavior before dissolution. Examples of this institution are divorce settlements (where judges might take into account the behavior of each party when dividing assets and custody of children) and labor arbitration hearings (where firms and workers agree to abide by an arbitrator’s decision).

The institution has the intuitive effect of increasing the outside option for cooperative players and decreasing it for defectors. Termination in this setting provides a more-powerful punishment, while use of the punishment is more palatable to those cooperative parties who feel certain their partners have been defecting.

Technically, the treatment induces a theoretically complex stochastic game with an imperfectly observed, endogenous state variable: the cooperation-difference $\omega_t = \sum_{k=1}^{t-1} \mathbf{1}\{a_1^k = C\} - \mathbf{1}\{a_2^k = C\}$. Each agent observes the history $\{a_i^k, Y^k\}_{k=1}^{t-1}$, from which they update their beliefs about ω_t . This belief then influences their expected payoff from taking the action $a_i^t = T$. For a patient enough player, individual rationality can be shown to be arbitrarily close to $\bar{\pi}$.³⁸

Rather than a technical analysis, our focus here will be on the behavior of human subjects in this environment. The environment is motivated by its intuitive punishment and the parallelism to extant institutions. However, we note that none of the simple pure-strategy machines (nor non-trivial mixed-strategy extensions) are SPE of the *A-Judge* game.

Rationality Hypothesis 8. *In the A-Judge treatment the discounted-average payoff vector satisfies $(W_1, W_2) \geq (\bar{\pi}, \bar{\pi})$.*³⁹

6. ASYMMETRIC RESULTS

For our asymmetric treatments we here focus on summarizing the main qualitative results. More-detailed tables, with analogues to figures and tables from Section 4, are included in the appendix for completeness. Instead, for this section we attempt to provide the reader with the main shifts in outcomes, drawing from the analysis techniques discussed in Section 4.

³⁸Each player can specify a strategy that cooperates initially. After every round it calculates the probability of the observed sequence of outcomes under the null that the other play is cooperating. The strategy terminates if the probability of the observed sequence drops below some pre-specified confidence level α^* . So long as α^* is small enough and δ large enough, ex ante this strategy guarantees an amount close to $\bar{\pi}$, regardless of the other player’s choices.

³⁹We present here the hypothesis for $\delta \rightarrow 1$. When $\delta = 4/5$ the individual rational payoff—if it exists, which is not guaranteed for games of this type—is strictly lower, and we can show it lies in the interval $[101.7, 111.4]$. Given arbitrary strategies σ_i and σ_j for the two players, and expected discounted-average payoff of $\hat{W}_i(\sigma_i, \sigma_j)$. The two bounds on $\min_{\sigma_i} \max_{\sigma_j} \hat{W}_i(\sigma_i, \sigma_j)$ come from: $101.7 = \min_{\sigma_j} W_i(\hat{\sigma}_i, \sigma_j)$ for the player strategy $\hat{\sigma}_i$ which plays two rounds of initial cooperation then terminates after the first failure; and $111.4 = \max_{\sigma_i} W_i(\sigma_i, \hat{\sigma}_j)$ for the strategy $\hat{\sigma}_j$ that randomizes between termination ($\frac{6}{11}$) and Always-Defect ($\frac{5}{11}$).

TABLE 5. Average Behavior in Asymmetric Treatments (Last 10 Supergames)

Treatment ($\bar{\pi}/\pi$)	Cooperation			Activity	Active Choices		N_S	N_T
	$t = 1$	$1 \leq t \leq 5$	$t = 5$	$2 \leq t \leq 5$	C	$T \mid \text{not } C$		
S-125 (\$1.25/\$1.25)	0.571	0.421	0.314	0.698	0.556	0.177	42	420
A-First (\$1.25/\$0.75)	0.089	0.024	0.007	0.011	0.117	0.968	27	270
A-Last (\$1.35/\$1.25)	0.479	0.368	0.286	0.755	0.458	0.085	29	290
A-Judge (\$1.25/\$0.75)	0.893	0.826	0.744	0.937	0.870	0.132	27	270

Note: Number of supergames are given for each session, supergame lengths are in terms of payment rounds, observed rounds are nearest multiple of five above the final payment round. Efficiency figures are for last ten supergames in each session.

Finding 6. *Results at the aggregate level indicate that*

- (A) *A division rewarding the termination party leads to a coordination on joint dissolution from the first opportunity and substantially lower efficiency.*
- (B) *A division rewarding the player being jilted leads to a cooperation rate similar to comparable symmetric treatments, but much lower use of termination as a punishment.*
- (C) *A division on termination that favors the more cooperative player leads to very high cooperation rates, higher efficiency, and lower inactivity.*

Where our symmetric treatments exhibit substantially similar cooperation rates, with greater differences in the strategies used to punish failures, our three asymmetric treatments demonstrate much greater variation in cooperation rates. Table 5 provides averages for the data from the last ten supergames from each asymmetric session, as well as the rates from the *S-125* treatment as a comparison.⁴⁰

Of the three asymmetric treatments, the clearest change from the symmetric results is in *A-First*. Subjects' average cooperation rate across the supergames is just 2.4 percent, which is far below the cooperation rates in our symmetric treatments (51 percent in *S-75* and 42 percent in *S-125* for the last ten supergames). Looking just at the first round behavior, the average cooperation rate across all of our symmetric treatments is 58.4 percent. Instead for *A-First* the first-round cooperation rate is 8.9 percent, with the vast majority of choices (88.9 percent) instead choosing to terminate at the very first opportunity. Using session-level averages this difference in first-round cooperation is highly significant ($p = 0.013$ for a Mann-Whitney test comparing the 11 symmetric sessions to the two for *A-First*). The high rate of termination leads to an overall activity rate of just 1.1 percent.

In efficiency terms, the high rate of termination and low joint payoff from its use leads to substantially reduced efficiency. Moreover, as joint termination is Pareto dominated by joint defection (our efficiency normalization) the *A-First* efficiency level is -45.6 percent. This is

⁴⁰Table 9 in the Appendix provides these same figures across all supergames.

lower than any of the symmetric treatments where the lowest treatment-level efficiency was 23.6 percent.⁴¹ Our results in this treatment point to a clear coordination on the worst-case outcome allowed by Rationality Hypothesis 6: joint termination in round one.

The second component of Finding 6 describes outcomes in *A-Last*, where the two dissolution payoffs have smaller difference than *A-First*, with the person initiating termination receiving a \$1.25 payoff, the other receiving \$1.35. Cooperation rates for *A-Last* are certainly closer to the levels found in our symmetric treatments, but lower than any of the symmetric termination treatments. This is true for the overall cooperation rates, as well as for just rounds one and five. However, the reduction in overall cooperation is not significantly different from *S-105-135* at the session level ($p = 0.154$).

While the lower cooperation levels for *A-Last* are not significantly different from similar symmetric treatments, we do observe significant differences in the rate at which subjects resort to termination as a punishment. Given non-cooperative behavior, subjects in the *A-Last* treatment are more likely to choose defection than termination in comparison to *S-105-135*. The two *A-Last* sessions are at an extreme compared to the six symmetric treatments, and we reject similar use of termination as a punishment using session-level averages ($p = 0.036$).

Drilling down to examine the conditional response in *A-Last* provides further detail (see Figure 5(B) in the appendix for the complete conditional response). For the *S-105-135* treatments termination is most commonly observed following failed cooperation, being chosen 27.2 percent of the time after this history. In contrast, the rate of termination in *A-Last* following failed cooperation is a third this at 9.2 percent.⁴² As cooperation following failed cooperation is higher by 6 percentage points in *S-105-135* than *A-Last*, the necessary choice shift is towards the only remaining option: defection. Where subjects in the *S-105-135* sessions choose to defect following failed cooperation at a rate of 25.2 percent, they do so at a 49.2 percent rate in *A-Last*. The small asymmetry introduced between the payoffs on termination leads to a clear shift towards in-relationship punishments.

The final component of Finding 6 concerns the *A-Judge* treatment, where the higher payoff ($\bar{\pi} = \$1.25$) is assigned to the more-cooperative partner and the lower payoff ($\underline{\pi} = \$0.75$) to the less-cooperative partner. We find the highest cooperation rates from any of our experimental manipulations in this treatment. Indeed, across all of the symmetric and asymmetric treatments, *A-Judge* is the only treatment where cooperation rates *increase* as the session proceeds.⁴³ The 89.3 percent rate of initial cooperation is significantly greater than the *S-75-135* treatments ($p = 0.013$ from a session-level Mann-Whitney test), and cooperation continues at this high rate across the supergame. By round five, close to three-quarters of subject choices are still cooperative, which is more than double the 36.7 percent

⁴¹The two *A-first* sessions are lower than all 11 symmetric sessions for cooperation, activity and payoff efficiency. As such all session-level Mann-Whitney tests will reject equivalence with the same p -value of 0.013.

⁴²While a probit (with subject-level clustering) fails to reject ($p = 0.598$) similar cooperation levels between *A-Last* and *S-105-135* conditional on failed cooperation the previous round, it strongly rejects equivalence over the probability of choosing termination after the same history ($p = 0.004$).

⁴³Across the first ten supergames the round-one cooperation rate is 81.9 percent, which increases to 89.3 percent in the last ten supergames. The increasing trend is significant ($p = 0.000$) according to a subject-level random-effect probit.

cooperation rate in the symmetric treatments at this point (significantly different with $p = 0.027$).

Looking at non-cooperative, active choices, termination is chosen 13.2 percent of the time. Focusing on behavior following failed cooperation, termination is chosen in 21.8 percent of these situations, with defection at 11.3 percent. The modal response to failed cooperation is lenience, to keep on cooperating. In the appendix we report strategy estimates for the asymmetric treatments alongside the symmetric treatments. For *A-Judge* the most-lenient option in the included strategy set (*Always C*) has a 60.3 percent estimated frequency. As the next-most-lenient strategy included within the estimation procedure allows for three failures, one interpretation of the 60.3 percent figure is that the majority of subjects in the *A-Judge* treatment require more than three failures before entering a punishment phase.⁴⁴ Looking at the estimated strategy types we find that 90.6 percent allow for ongoing cooperation, which is significantly greater than our next most-cooperative treatment group (*S-75-95*, $p = 0.011$).⁴⁵ Regardless of the method of analysis, the *A-Judge* division rule leads to a clear coordination on cooperative outcomes.

Examining the efficiency within the *A-Judge* treatment in the last ten supergames we find a figure of 76.8 percent. Though this is higher than any of our symmetric treatments ($p = 0.027$), given the high cooperation rates, readers might expect this figure to be higher, closer to the fully-cooperative upper bound. The reason the efficiency figure is not higher is that though the termination punishments are rarely used—activity is over 90 percent—when it is chosen termination is highly inefficient. As such, the large efficiency gains from greater cooperation over a high-outside-option treatment (for example, *S-125*) are lowered due to the very inefficient outcome when dissolution occurs.

Asymmetric Treatment Summary. Our asymmetric treatments change the payoffs on termination, so that one party gets a larger amount, where the assignment is a function of behavior within the relationship. We find that changes to the asymmetric division rule lead to substantial shifts in subjects' behavior within the relationships.

Where a focus in the relational contracting literature has been on efficient *ex post* allocation after a relationship ends—for example Texas shoot-outs, and the pricing and division of assets—there has been less work examining how such divisions affect outcomes *ex ante*. Our symmetric experimental results suggest that that inefficient division on dissolution does in itself guarantee more-cooperative outcomes. In addition, our *A-First* treatment has subjects coordinate on inefficient dissolution, with almost all relationships terminated at the first opportunity. However, our similarly inefficient *A-Judge* division rule does produce a clear coordination on efficient play.⁴⁶ Our asymmetric treatments illustrate the large effects that differences in the division rule on dissolution can have on in-relationship behavior. That is,

⁴⁴One problem with the SFEM estimation procedure in *A-Judge* is that the heavy coordination on cooperation leads to lower power, as less punishment phases are entered, particularly those with long sequences of failures. For example, in the last ten supergames only five of the 270 supergames have failed cooperation in two *consecutive* rounds, and none of them have failed cooperation in three consecutive rounds.

⁴⁵In terms of punishments, the estimates indicate that 20.0 percent of strategies use termination at some point, 13.3 percent are forgiving (73.6 percent if we include *Always C*), and 68.8 percent are lenient.

⁴⁶Though we do not observe a substantial direct effect on cooperation in *A-Last*, the treatment does indicate how small changes to the division on termination can substantially alter punishments.

it is not the total inefficiency on dissolution that drives better outcomes, it is instead the way in which it is distributed.

While our asymmetric experiments can not hope to be a comprehensive exercise—given the vast possibilities for division rules—the three treatments above do distill strategic tensions plausibly present outside the laboratory. Given that folk theorems for ongoing relationships will always lead to indeterminacy in prediction without strong selection assumptions, our asymmetric treatments show how the laboratory can provide a test-bed for examining how designed features of a contract can affect behavior in long-run relationships.

7. CONCLUSION

We experimentally investigate a series of prisoner’s dilemma games with imperfect monitoring. Introducing a termination option into the PD stage game that can unilaterally end the relationship, our experiments manipulate the value of the outside option available to the players.

Our first set of symmetric treatments examine the effect from varying a commonly available outside option, where each partner receives the same payoff on dissolution. Here we contrast observed outcomes to the same imperfect-monitoring environment without a termination option. Our findings suggest that the use of termination as a punishment is directly related to the outside option’s value: if remaining in an uncooperative relationship stochastically dominates walking away subjects rarely end the relationship. However, while the *presence* of a dissolution option does increase both cooperation and the lenience of the dynamic strategies used, we do not observe differences in the cooperation rate with respect to the *value* of the outside option.

One slight difference between behavior and theory in our experiments is that subjects begin to use termination at more substantial frequencies once the outside option’s value exceeds the lowest individually rational realization, as opposed to its expectation. Though termination is not used to end the relationship when it is too costly (stochastically dominated), subjects do make use of it at non-negligible rates in treatments where its expected value is below that guaranteed within the relationship. This leads to a non-monotonicity for efficiency with respect to the outside option. Though somewhat constant for lower values, there is a sharp drop once termination begins to be used as a punishment. Subsequent increases to the outside option then leads to increasing efficiency, as they reduce the costliness of the punishment. The results suggest that symmetric frictions in the dissolution process (incompleteness in contracts, costs for litigation, etc.) can decrease efficiency. Such costs are ideally either: (i) minimized; or (ii) small enough to make dissolution behaviorally plausible, but large enough so that the vast majority are unwilling to use this nuclear option.

Our second set of treatments illustrate another direction: changing the distribution of payoffs on termination. In our asymmetric treatments one partner gets a higher payoff, where the assignment of the better outcome is dependent on the choices made while the partnership was active. Three institutions are analyzed: (i) the party that terminates gets a higher payoff; (ii) the party being terminated gets a higher payoff; and, finally, (iii) the party that was more cooperative gets the higher payoff. In efficiency terms, changes to the distribution of payoffs have much larger effect sizes than observed in the symmetric outside-option treatments.

Where first-movers receive the better termination payoff we see almost all subjects coordinate on termination in the very first round. Moreover, the precise dissolution payoffs in this treatment were chosen so that joint-termination is Pareto dominated by the in-relationship option. Despite subjects having access to the same in-relationship payoffs as our other treatments, this treatment has starkly lower efficiency than any of our other experimental treatments. The treatment underlines the ways in which poorly chosen institutional or contractual features can have highly detrimental effects on outcomes.

In our second asymmetric treatment where the party being terminated receives the higher payoff, we observe a large drop off in the fraction of subjects using termination, relative to comparable symmetric treatments. Though the small asymmetry in outcomes on termination in this treatment does not lead to a substantial shift in cooperation rates relative to comparable symmetric treatments, it does cause subjects to reduce their use of termination. Instead, subjects switch to in-relationship punishments, where the treatment demonstrates how small asymmetries can alter the selection of on-path punishments.

Our final asymmetric payoff treatment, *A-Judge*, is the most successful across our treatments, where selected strategies are highly cooperative and efficient. This treatment determines who receives the higher and lower dissolution payoffs by examining who cooperated the most within the relationship while it was active. The treatment is an attempt to model an arbitrator or judge choosing how to distribute payoffs through a moral remit to assign it to the “more-deserving” party.⁴⁷ Though the termination punishment in our parametrization is very inefficient when used, the rule determining the division between the partners allows for a harsher punishment for those deviating from the cooperative path, and can be a rational choice for cooperators. These features of the punishment lead to very high cooperation rates, which more than compensates for the inefficiencies generated when termination is used along the path.

Where our symmetric treatments indicate a great deal of regularity in the rate of cooperation, with selected punishments dictated by individual rationality, our asymmetric treatments instead show a large variation in the chosen behaviors. Rather than being a comprehensive exercise, our asymmetric results serve to illustrate the degree to which different division rules on dissolution can cause relationships to take very different paths. Methodologically, our asymmetric treatments provide clear examples of the power of the laboratory as a tool to understand the *total* effects of contractual features on long-run relationships. As a more-specific result, our asymmetric treatments make clear that while inefficiency in the dissolution process can be useful for increasing ex ante efficiency, it does so only if the ex post costs of dissolution can be concentrated on poorly performing partners.

REFERENCES

- Abreu, D., D. Pearce, and E. Stacchetti, “Optimal cartel equilibria with imperfect monitoring,” *Journal of Economic Theory*, 1986, 39 (1), 251–269.
- , ———, and ———, “Toward a theory of discounted repeated games with imperfect monitoring,” *Econometrica*, 1990, pp. 1041–1063.

⁴⁷In our experiments the judge is automated and has access to a perfect forensic process. Future research might examine the extent to which the highly efficient outcomes we observe are retained when the arbitrator is a human subject without a set division rule and/or with imperfect information on the two player’s actions.

- Aoyagi, M. and G. Fréchet**, “Collusion as public monitoring becomes noisy: Experimental evidence,” *Journal of Economic Theory*, 2009, 144 (3), 1135–1165.
- Blonski, Matthias and Giancarlo Spagnolo**, “Prisoners’ other dilemma,” *JW Goethe University, Frankfurt, Economics Working Paper*, 2004.
- Bó, P. Dal**, “Cooperation under the shadow of the future: experimental evidence from infinitely repeated games,” *The American Economic Review*, 2005, 95 (5), 1591–1604.
- and **G.R. Fréchet**, “The evolution of cooperation in infinitely repeated games: Experimental evidence,” *American Economic Review*, 2007, 50, 54.
- Bó, Pedro Dal and Guillaume R Fréchet**, “The evolution of cooperation in infinitely repeated games: Experimental evidence,” *American Economic Review*, 2011, 101 (1), 411–429.
- and ———, “On the Determinants of Cooperation in Infinitely Repeated Games: A Survey,” *Journal of Economic Literature*, forthcoming.
- Brooks, Richard RW, Claudia M Landeo, and Kathryn E Spier**, “Trigger happy or gun shy? Dissolving common-value partnerships with Texas shootouts,” *RAND Journal of Economics*, 2010, 41 (4), 649–673.
- Comino, Stefano, Antonio Nicolò, and Piero Tedeschi**, “Termination clauses in partnerships,” *European Economic Review*, 2010, 54 (5), 718–732.
- Cramton, Peter, Robert Gibbons, and Paul Klemperer**, “Dissolving a Partnership Efficiently,” *Econometrica*, 1987, 55 (3), 615–632.
- Dreber, A., D.G. Rand, D. Fudenberg, and M.A. Nowak**, “Winners don’t punish,” *Nature*, 2008, 452 (7185), 348–351.
- Duffy, J. and J. Ochs**, “Cooperative behavior and the frequency of social interaction,” *Games and Economic Behavior*, 2009, 66 (2), 785–812.
- Dutta, Prajit K**, “A folk theorem for stochastic games,” *Journal of Economic Theory*, 1995, 66 (1), 1–32.
- Embrey, Matthew, Guillaume Fréchet, and Ennio Stacchetti**, “An Experimental Study of Imperfect Public Monitoring: Efficiency versus Renegotiation-Proofness,” 2013.
- Fréchet, Guillaume R and Sevgi Yuksel**, “Infinitely Repeated Games in the Laboratory: Four Perspectives on Discounting and Random Termination,” February 2013. NYU working paper.
- Friedman, J.W.**, “A non-cooperative equilibrium for supergames,” *Review of Economic Studies*, 1971, 38 (1), 1–12.
- Fudenberg, D. and E. Maskin**, “The folk theorem in repeated games with discounting or with incomplete information,” *Econometrica*, 1986, pp. 533–554.
- , **D. Levine, and E. Maskin**, “The folk theorem with imperfect public information,” *Econometrica*, 1994, pp. 997–1039.
- Fudenberg, Drew and Jean Tirole**, *Game Theory* 1991.
- , **David G Rand, and Anna Dreber**, “Slow to anger and fast to forgive: Cooperation in an uncertain world,” *The American Economic Review*, 2012, 102 (2), 720–749.
- Gaudeul, Alexia, Paolo Crosetto, and Gerhard Riener**, “On the stability of partnerships when individuals have outside options, or why allowing exit is inefficient,” 2015.
- Green, E.J. and R.H. Porter**, “Noncooperative collusion under imperfect price information,” *Econometrica*, 1984, pp. 87–100.

- Hauk, Esther and Rosemarie Nagel**, “Choice of Partners in Multiple Two-Person Prisoner’s Dilemma Games An Experimental Study,” *Journal of Conflict Resolution*, 2001, 45 (6), 770–793.
- Honhon, Dorothee and Kyle B Hyndman**, “Flexibility and Reputation in Repeated Prisoners’ Dilemma Games,” 2015. University of Texas working paper.
- Hyndman, Kyle and Dorothee Honhon**, “Flexibility in Long-Term Relationships: An Experimental Study,” 2014. University of Texas working paper.
- Li, J. and E. Wolfstetter**, “Partnership dissolution, complementarity, and investment incentives,” *Oxford Economic Papers*, 2010, 62 (3), 529–552.
- Maskin, Eric and Jean Tirole**, “Markov perfect equilibrium: I. Observable actions,” *Journal of Economic Theory*, 2001, 100 (2), 191–219.
- McAfee, R Preston**, “Amicable divorce: Dissolving a partnership with simple mechanisms,” *Journal of Economic Theory*, 1992, 56 (2), 266–293.
- Murnighan, J.K. and A.E. Roth**, “Expecting continued play in prisoner’s dilemma games,” *Journal of Conflict Resolution*, 1983, 27 (2), 279–300.
- Palfrey, T.R. and H. Rosenthal**, “Repeated play, cooperation and coordination: An experimental study,” *Review of Economic Studies*, 1994, pp. 545–565.
- Porter, R.H.**, “Optimal cartel trigger price strategies,” *Journal of Economic Theory*, 1983, 29 (2), 313–338.
- Radner, R., R. Myerson, and E. Maskin**, “An example of a repeated partnership game with discounting and with uniformly inefficient equilibria,” *Review of Economic Studies*, 1986, 53 (1), 59–69.
- Rand, D.G., S. Arbesman, and N.A. Christakis**, “Dynamic social networks promote cooperation in experiments with humans,” *Proceedings of the National Academy of Sciences*, 2011, 108 (48), 19193–19198.
- Roth, A.E. and J.K. Murnighan**, “Equilibrium behavior and repeated play of the prisoner’s dilemma,” *Journal of Mathematical Psychology*, 1978, 17 (2), 189–198.
- Vespa, Emanuel and Alistair J Wilson**, “Experimenting with equilibrium selection in dynamic games,” 2015. University of Pittsburgh working paper.

APPENDIX A. STRATEGY FREQUENCY ESTIMATION (FOR ONLINE PUBLICATION)

We will briefly describe the econometric model adopted in SFEM, discuss about the set of strategies we include in the estimation and then reports the estimated frequency of strategies for each of the treatments. To use SFEM, we assume that each subject chooses a fixed strategy for the last ten supergames. These chosen strategies are implemented with the possibility of independent mistakes, where another choice than the intended action is made. In this section, we report the results from our SFEM estimations. We first briefly describe the econometric model adopted and then report the estimated strategy weights for each of the main treatments.

Denote the choice made by subject i in round t of supergame m by c_{imt} . We specify a priori a set of $K = 38$ possible public perfect strategies $\Phi = \{\phi_1, \dots, \phi_K\}$, where the choice of the strategy ϕ_k prescribes the choice $s_{imt}^k = \phi(y_{jm1}, \dots, y_{jm(t-1)})$ following the public history $(y_{jm1}, \dots, y_{jm(t-1)})$. Each subject is assumed to follow a particular strategy ϕ_k , but they make independent mistakes each round with a probability $(1 - \beta)$, and chose the prescribed action with probability β . Given three actions (C , D and T) a random uniform choice is represented by $\beta = 1/3$, while a perfect match with a strategy by $\beta = 1$. The econometric model assumes a mixture model across the available strategies, so that strategy ϕ_k is selected with probability q_k .

Define the indicator, $I_{imt}^k = 1\{c_{imt} = s_{imt}^k\}$, which assigns a value of one if the observed subject choice matches the strategy choice σ_k match. The likelihood that the observed choices for subject i were generated by the strategy ϕ_k are given by

$$\Pr_i(\phi_k; \beta) := \prod_{m \in M_i} \prod_{t=1}^{T_{im}} \beta^{I_{imt}^k} (1 - \beta)^{1 - I_{imt}^k},$$

where M_i is the set of supergames, and T_{im} the set of active rounds. Combining across all subjects in a treatment we obtain the following likelihood function:

$$\sum_{i \in \mathcal{I}} \ln \left(\sum_{\phi_k \in \Phi} q_k \Pr_i(\phi_k; \beta) \right)$$

for the specified set of strategies Φ and summing the log-likelihoods across all subjects \mathcal{I} in the treatment. The parameters to be estimated by maximum likelihood are the vector of probabilities $\mathbf{q} = (q_1, \dots, q_K)$ and the strategy-match probability β , under constraints that $\beta \in [1/3, 1]$, and that the vector of probabilities \mathbf{q} lies in the probability K -simplex. The numerical maximization, and bootstrapping of the results, were completed in *Mathematica* using a differential-evolution constrained-optimization algorithm, using starting points for the estimation obtained following the same techniques followed in Dal Bó and Fréchette (2011).⁴⁸

In total we allow for 38 different strategies, motivated by the previous experimental literature (in particular Embrey et al., 2013, which has the closest signal space to our own study). The precise strategies used are described in Table 6. Table 7 presents the results from the estimation procedure, where we have left blank all estimates with a zero weight to

⁴⁸We also conducted the same exercise in *Matlab* using modified code provided by Guillaume Fréchette, and obtained qualitatively similar results. However, the *Mathematica* numerical routines seemed to be better at attaining a global solution.

make the strategies with positive estimates clearer. However, readers should note that all 38 strategies in Φ were included in each estimation.

TABLE 6. Strategy Descriptions

Strategy	Abbreviation	Description
Always Cooperate	AC	Always choose C
Always Defect	AD	Always choose D
Always Terminate	AT	Always choose T
Alternator (C)	CDCD	Start by cooperating; alternate between cooperation and defection
Alternator (D)	DCDC	Start by defecting; alternate between cooperation and defection
False Cooperator	C-AllD	Choose to cooperate in the first round; play defect in all subsequent rounds
One round cooperator	C-T	Cooperate in first round, then terminate in all subsequent rounds
One round defector	D-T	Defect in first round, then terminate in all subsequent rounds
Grim trigger	Grim	Cooperate until first observed failure, defect thereafter
Monotone	Mono	Cooperate initially, then cooperate after successes, defect after failures
Win-stay-lose-shift	WSLS	Cooperate initially, repeat last action a_j on a success, and $\{C, D\} \setminus \{a_j\}$ otherwise
One-Period Punishment	T11 CD	Cooperate initially. Continue to cooperate after successes; after a failure defect for one period and return to cooperating the next round regardless of the outcome.
Suspicious Monotone	S-Mono	Same as Mono, but start by defecting.
	S-WSLS	Same as WSLS, but start by defecting.
Lenient Grim k	Grim- k	Cooperate until k observed total failures, defect thereafter. Variants included are Grim-2 and Grim-3.
Monotone i -Cooperate j -Defect	Mono- ij	Machine states \mathcal{A}_1 to \mathcal{A}_{i+j} , associated with the action C for states \mathcal{A}_1 to \mathcal{A}_i and D for \mathcal{A}_{i+1} to \mathcal{A}_{i+j} . Start in state \mathcal{A}_1 . In any round t in state \mathcal{A}_j success moves the machine to state \mathcal{A}_{j-1} (\mathcal{A}_1 at lower limit) and failure moves it to the state \mathcal{A}_{j+1} (\mathcal{A}_{i+j} at top limit). Included variants are Mono-21, Mono-31, Mono-12, Mono-22.
Sum of success	Sum- k	Cooperate if there were at least as many successes as failures in the last k rounds. Variants included are Sum2, Sum3 and Sum4
Suspicious sum of success	S-Sum2	Start out defecting, henceforth cooperate if one or more success in last 2 rounds
Cooperate, k -strikes you're out	C- k -Strike	Cooperate until k observed failures, terminate thereafter. Variant included are C-1-Strike, C-2-Strike and C-3-Strike
Defect, k -strikes you're out	D- k -Strike	Defect until k observed failures, terminate thereafter. Variant included are D-1-Strike, D-2-Strike and D-3-Strike
Demotion- ij strike out	CD- ij -Strike	Machine states \mathcal{A}_1 to \mathcal{A}_{i+j+1} , associated with the action C for states \mathcal{A}_1 to \mathcal{A}_i , D for \mathcal{A}_{i+1} to \mathcal{A}_{i+j} , and T for \mathcal{A}_{i+j+1} . Start in state \mathcal{A}_1 . In any round t in state \mathcal{A}_k keeps the machine in state \mathcal{A}_k , failure moves the machine to the state \mathcal{A}_{k+1} (\mathcal{A}_{i+j+1}). Included variants are CD-11-Strike, CD-12-Strike, CD-21-Strike and CD-22-Strike.
Probation period ij	Probation- ij	Machine states \mathcal{A}_1 to \mathcal{A}_{i+j+1} , associated with the action C for states \mathcal{A}_1 to \mathcal{A}_i , D for \mathcal{A}_{i+1} to \mathcal{A}_{i+j} , and T for \mathcal{A}_{i+j+1} . Start in state \mathcal{A}_1 . In any round t in state \mathcal{A}_k moves the machine to state \mathcal{A}_{k-1} (\mathcal{A}_1 at lower limit), failure moves the machine to the state \mathcal{A}_{k+1} (\mathcal{A}_{i+j+1} which is absorbing). Included variants are Probation (i.e Probation-11), S-Probation (a variant of Probation-11 that starts in state \mathcal{A}_2), Probation21 and Probation22

TABLE 7. Strategy Estimates: Main Treatments (last ten supergames)

	None	Sym-75-95	Sym-105-135	Asym-First	Asym-Last	Asym-Moral
AC	0.027 (0.066)	0.184* (0.103)	0.073** (0.036)		0.218** (0.103)	0.603** (0.252)
AD	0.390*** (0.145)	0.325*** (0.098)	0.240*** (0.065)		0.375*** (0.136)	0.020 (0.025)
AT			0.056** (0.025)	0.932*** (0.077)	0.036 (0.034)	
CDCD			0.031* (0.018)	0.068 (0.050)		0.037 (0.031)
DCDC	0.021 (0.024)		0.014 (0.012)			
C-AllD		0.017 (0.013)	0.053** (0.025)		0.033 (0.033)	0.037 (0.050)
C-T						
D-T						
Grim	0.247*** (0.093)	0.134*** (0.039)			0.202*** (0.075)	0.010 (0.008)
Mono	0.064** (0.031)	0.056** (0.022)	0.072*** (0.018)		0.005 (0.040)	
WSLS		0.015 (0.015)	0.012 (0.014)			0.036 (0.044)
T11 CD		0.017 (0.014)				
S-Mono	0.047** (0.019)					
S-WSLS	0.023 (0.020)		0.011 (0.011)			0.018 (0.026)
Grim-2		0.015 (0.010)			0.040 (0.030)	
Grim-3	0.031 (0.026)	0.010 (0.011)				
Mono21		0.026 (0.017)	0.011 (0.016)		0.021 (0.036)	
Mono31			0.062* (0.036)			
Mono12	0.106*** (0.037)	0.048** (0.022)				0.003 (0.011)
Mono22		0.020* (0.010)	0.050** (0.021)			

	None	Sym-75–95	Sym-105–135	Asym-First	Asym-Last	Asym-Moral
Sum2	0.045*** (0.012)	0.024** (0.010)				
Sum3		0.035*** (0.010)				
Sum4						
S-Sum2		0.021** (0.010)	0.032** (0.016)			0.037 (0.042)
C-1-Strike		0.013 (0.013)	0.066* (0.036)		0.024 (0.027)	0.059 (0.078)
C-2-Strike		0.017 (0.017)	0.095** (0.037)			0.048 (0.090)
C-3-Strike		0.007 (0.015)	0.056** (0.027)			
D-1-Strike			0.022 (0.014)		0.031 (0.029)	
D-2-Strike			0.031 (0.020)		0.015 (0.025)	
D-3-Strike						
CD-11-Strike						
CD-12-Strike						
CD-21-Strike						0.054 (0.067)
CD-22-Strike						
Probation		0.015 (0.010)				0.039** (0.017)
S-Probation			0.011 (0.012)			
Probation21						
Probation12						
			(0.010)			
β	0.830 (0.033)	0.899 (0.021)	0.860 (0.017)	0.892 (0.053)	0.814 (0.028)	0.908 (0.022)

Note: Bootstrapped standard errors (across sessions, subjects and supergames) in parentheses. Significance indicated by: ***–1 percent level; **–5 percent level; *–10 percent level. Strategies with zero-weight estimates are not included in the table cells, note though that most of these zero-weight coefficients have non-zero realizations for the bootstrapped standard errors.

APPENDIX B. SUPPLEMENTAL FIGURES AND INSTRUCTIONS (FOR ONLINE PUBLICATION)

TABLE 8. Asymmetric Treatment Strategy Classifications (Last ten supergames)

	<i>A-First</i>	<i>A-Last</i>	<i>A-Judge</i>
Initial <i>C</i>	6.8%	54.2%	92.6%
Ongoing <i>C</i>	0.0%	50.9%	90.6%
Lenient	0.0%	27.9%	74.2%
Forgiving	0.0%	2.6%	13.3%
(+ <i>All-C</i>)	0.0%	24.4%	73.6%
Terminating	93.2%	10.7%	20.0%

TABLE 9. Average Behavior in Asymmetric Treatments (All Supergames)

Treatment	Cooperation			Activity $2 \leq t \leq 5$	Active Choices		N_S	N_T
	$t = 1$	$1 \leq t \leq 5$	$t = 5$		<i>C</i>	<i>T</i> not <i>C</i>		
S-125	0.602	0.472	0.368	0.758	0.586	0.153	42	774
A-First	0.243	0.101	0.047	0.101	0.360	0.827	27	465
A-Last	0.578	0.458	0.336	0.832	0.529	0.080	29	650
A-Judge	0.844	0.770	0.683	0.919	0.824	0.140	27	463

Note: Number of supergames are given for each session, supergame lengths are in terms of payment rounds, observed rounds are nearest multiple of five above the final payment round. Efficiency figures are for last ten supergames in each session.




FIGURE 5. Conditional response in last ten supergames, $a_i^t \mid (a_i^{t-1}, y^{t-1})$
Note: For each treatment the horizontal axis indicate the proportion of active rounds $t \geq 2$ in the last five supergames given the relevant (a_i^{t-1}, y^{t-1}) history. The vertical axis indicates the proportion of each of the three possible actions: D in gray, C in white, and the *Terminate* action in black. The bar on the left of each plot provides the proportion of actions chosen at the empty history (round 1)

MATCH 1 ROUND 4

END-PARTNERSHIP PAYOFF = 125

Remaining Time: 30

**OUTCOME
PAYOFF**

Success	200
Failure	50

ACTION PAYOFF

A	-50
B	50

CHANCE OF SUCCESS

Your/Your Partner's Choices	A	B
A	99%	50%
B	50%	10%

Partner Red

A B
END PARTNERSHIP

Partner Blue

A B
END PARTNERSHIP

OK

Round	Your Action	Outcome	Payoff	Round	Your Action	Outcome	Payoff
3	A	Failure	0	3	B	Success	250
2	A	Success	150	2	B	Success	250
1	A	Success	150	1	B	Success	250

FIGURE 6. Screenshot of Experimental Interface

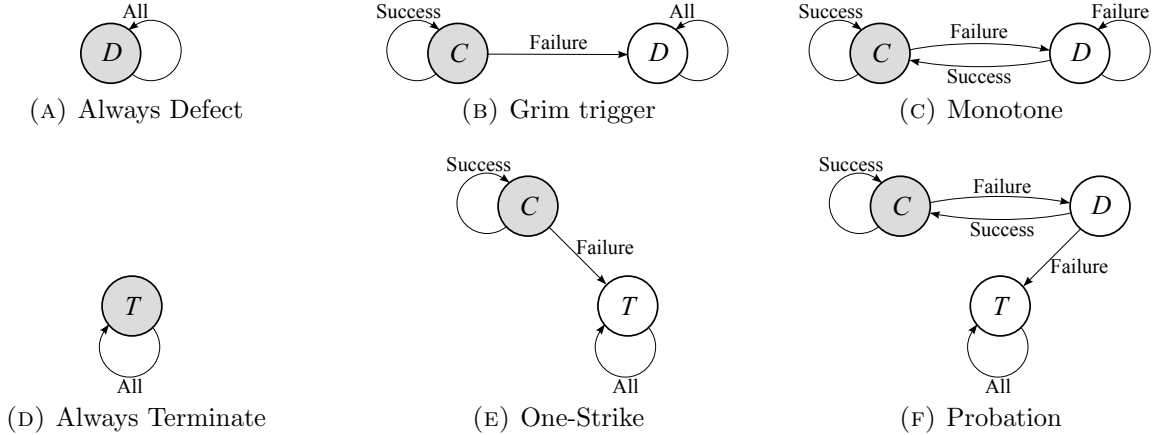


FIGURE 7. Simple three-state machines with Termination

APPENDIX C. EQUILIBRIUM STRATEGIES

In this appendix we outline simple equilibria of our games with $\delta = 4/5$. Our focus will be on the following restricted, but cognitively simple set of equilibria: weakly undominated, symmetric, stationary, public-perfect equilibria (wSSPPE), which can be represented as finite machines. That is, we will look for strategies of the game that depend only on the public signal in the last round Y^{t-1} and a finite number of internal states. The machine's chosen action at each point in time is dictated by its internal state, and the transition between these states depends only on the public signal (\mathcal{S} or \mathcal{F}) and the internal state last round. For simplicity, we examine machines with three internal states C , D and T , where the three states are connected to the actions cooperate, defect, and terminate, respectively. By discussing only the strategies that fall into this small category we hope to illustrate the broad effects on the types of equilibria possible from changes to the outside option. In our data analysis, we will consider a larger constellation of machines, and asymmetric pairings between them.

Because termination is an absorbing state by construction, the set of 3-state machines in our setting has 81 distinct entries.⁴⁹ In terms of theoretical prediction, from those 81 possible machines, just three form wSSPPEs when the symmetric outside option π is less than \$1.15 (or when dissolution is not present). These three machines are depicted in the Figure 7, and correspond to the strategies (A) *Always Defect*, (B) the *Grim Trigger*, and (C) the *Mono(tone)* strategy, where we have highlighted the most-efficient starting state in gray.⁵⁰ When the value of termination is strictly greater than \$1.15 the set of equilibria changes, *Always Defect* is now replaced by the strategy (D) *Always Terminate* as an equilibrium, while the *Grim Trigger* is replaced by its termination complement (E), which we will refer

⁴⁹The C and D states each have two edges, with three possible destinations, so there are $3^4 = 81$ possible machines, where the machines can start in any state

⁵⁰If we look at the set of all two-state machines without termination (allowing for asymmetric matchings between different machines) there are four equilibria (unique up to state relabeling): symmetric *Always Defect*, symmetric *Mono* starting in either the *Cooperate* or *Defect* state, and symmetric *Grim Trigger*.

to as *1-strike*.⁵¹ However, the forgiving *Mono* machine remains a symmetric equilibrium so long as the value of termination is not too much greater than \$1.15. As the termination value increases past \$1.24, *Mono* stops being an equilibrium—the continuation value in the *D* state falls below the dissolution option π , so the punishment is no longer sub-game perfect. As the value of termination increases further still (beyond \$1.31) the *1-strike* strategy drops from the equilibrium set too, as defecting in the *C* state becomes profitable as the dissolution punishment lacks power.^{52,53}

Other than *Mono*, of the 81 machines we consider there are no other forgiving wSSPPEs for any termination value π —strategies capable of returning to cooperation after entering a punishment phase.⁵⁴ The only forgiving machine with incentive-compatible cooperation using both dissolution and the *D* action in the punishment path is the *Probation* strategy illustrated in Figure 7(F).⁵⁵ However, despite the punishment supporting cooperation, the strategy is not a wSSPPE as the punishment phase is not incentive compatible. Best-responding agents will deviate to play *C* (or *T* if π is large enough) where the strategy specifies the *D* action. This is because of a much higher probability of returning to the high-payoff cooperation state if they deviate to *C* (a 50 percent chance) as opposed to *D* (a 10 percent chance). Other simple forgiving strategies such as *Win-stay-lose-shift* (WSLS) and the unconditional one-round-punishment (*T11*) lack incentive compatibility for cooperation, as continuation values in the in-relationship punishment are too high.⁵⁶

The asymmetric-division treatments have similarly stark predictions. For *A-First* just two of the three-state machines are wSSPPE: *Always Terminate* and *1-Strike*. In the *A-Last* treatment, none of these 81 machines are symmetric equilibria, though asymmetric combinations such as *Always Defect/Always Terminate* and *Grim/1-Strike* are w(no S)SPPEs. Finally, in the *A-Judge* treatment, none of the 81 machines, nor any asymmetric pairing are

⁵¹A suspicious version of *1-strike* that starts out defecting (transiting to the cooperation phase of *1-strike* either unconditionally or coordinated on either signal realization) is also an admissible equilibrium for $115 \leq \pi \leq 131.25$.

⁵²Moreover, once the value of termination crosses \$1.32, the *only* perfect equilibrium of any form for $\delta = 4/5$ is *Always Terminate*.

⁵³In addition to the existence of differing equilibria, when termination increases in value past \$1.15, the comparable risk-dominance orderings of the equilibria change. Introduced in Blonski and Spagnolo (2004) and applied by Dal Bó and Fréchette (2011), risk dominance has been a useful measure to predict subjects' response in perfect monitoring environments. Among the three equilibrium strategies for $\pi_T < \$1.15$ (*Grim*, *Mono* and *Always Defect*), we find that *Always Defect* risk dominates *Grim*, which in turn risk dominates *Mono*. Where $\pi_T \geq \$1.15$, *Mono* and *1-strike* both risk dominate *Always Terminate*, which dominates *Always Defect*.

⁵⁴No “lenient” strategies are possible without allowing for additional states that play the *C* action.

⁵⁵A variant of the *Probation* strategy which exchanges the *Success* and *Failure* arrows in the *D*-state is also not an equilibrium, though here because cooperation is not incentive compatible when the termination value is weakly greater than \$1.15, and for lower values the termination action is weakly dominated.

⁵⁶There are many asymmetric equilibria when $\pi = \$1.15$, involving combinations of *Defection* and *Termination*. Once the termination value increases to \$1.25 *all* conditionally cooperative asymmetric equilibria combining the 81 machines involve at least one player using a variant of *1-strike*. For example, one player use the *Suspicious 1-strike* that starts at *Defect*, moves to *Terminate* on a failure, and to the standard *1-Strike Cooperation* state on a success; and the other player can use *Suspicious-Grim* (replace the *T* action in *Suspicious 1-strike* with *Always D*).

PPEs of the game. The *A-Judge* game requires much more sophisticated asymmetric strategies to form an equilibrium outcome (in particular mixed strategies with a corresponding belief update rule over the game state ω_t).