

A Continuous Dilemma[†]

By DANIEL FRIEDMAN AND RYAN OPREA*

We study prisoners' dilemmas played in continuous time with flow payoffs accumulated over 60 seconds. In most cases, the median rate of mutual cooperation is about 90 percent. Control sessions with repeated matchings over eight subperiods achieve less than half as much cooperation, and cooperation rates approach zero in one-shot sessions. In follow-up sessions with a variable number of subperiods, cooperation rates increase nearly linearly as the grid size decreases, and, with one-second subperiods, they approach continuous levels. Our data support a strand of theory that explains how capacity to respond rapidly stabilizes cooperation and destabilizes defection in the prisoner's dilemma. (JEL C72, C78, C91)

The computerized ATPCO system, introduced in the early 1990s, allowed airlines to rapidly adjust prices and to monitor those of rivals. The result, according to the US Department of Justice, was anticompetitive behavior—firms were able to more easily cooperate to keep prices high, costing consumers up to 2 billion dollars (Borenstein 2003, Klein 1998).

Does the ability to rapidly adjust actions actually encourage cooperation? Compared to the one-shot or discrete time strategic interactions usually analyzed by game theorists, are outcomes much different when choices are made asynchronously, in continuous time? These questions are not merely of theoretical interest because many modern interactions, ranging from just-in-time team production to e-commerce pricing, involve asynchronous strategic decisions made in real time.

In this article we take such questions to the laboratory. In a continuous time setting, we study variants of the prisoner's dilemma, the simplest and most famous example of a strategic tension between efficient cooperation and inefficient self-interest. Pairs of laboratory subjects are matched anonymously in 60-second periods, within which they can switch freely between cooperation and defection. They accrue flow payouts from one of four parametric variants of the prisoner's dilemma and then are randomly rematched for the next period. Each session runs 32–36 periods. We also

*Friedman: Economics Department, University of California, Santa Cruz, 1156 High Street, Santa Cruz, CA 95064 (e-mail: dan@ucsc.edu); Oprea: Economics Department, University of California, Santa Cruz, 1156 High Street, Santa Cruz, CA 95064 (e-mail: roprea@ucsc.edu). We are grateful to the National Science Foundation for support under grant SES-0925039, and to seminar audiences at Berkeley, Cal Poly SLO, Chapman, Harvard, NYU, Ohio State, Oxford, Santa Clara, and UCSC. We thank Stefano Della Vigna, Guillaume Frechette, Drew Fudenberg, Paul J. Healy, Paul Heidhues, John Kagel, Botond Kőszegi, David Laibson, Dan Levin, James Peck, Leo Simon, Ennio Stacchetti, and Michael Wellman for helpful comments, James Pettit for crucial programming support, Keith Henwood for excellent research assistance, and three anonymous referees and especially the editor for painstaking guidance in revising earlier versions.

[†] To view additional materials, visit the article page at <http://dx.doi.org/10.1257/aer.102.1.337>.

run control sessions with one-shot periods and with repeated discrete time subperiods, using identical payoff matrices, period lengths, and matching procedures.

Section I recalls some previous theoretical and experimental work related to our investigation, and Section II lays out our experimental design. Section III reports first results: continuous time enables median mutual cooperation rates of 90 percent or more in most of the variants, more than double the rates seen in our discrete (eight stages per period) repeated games, while mutual cooperation becomes quite rare in the one-shot sessions. The underlying strategies seem broadly consistent with some of the theoretical literature, particularly with Radner (1986), and Simon and Stinchcombe (1989).

Adapting previous theory to a continuous time setting with fast but not instantaneous reactions, Section IV obtains predictions arising from a class of epsilon equilibria in cutoff strategies. It then reports a second wave of sessions that vary the number of discrete stages from two to 60 within each 60-second period. As predicted, the data show a negative, almost linear, relationship between the cooperation rate and the length of the stage game. Indeed, the mutual cooperation rates with two stages per period are not far from zero, and those with 60 stages are not far from the rates seen in continuous time.

Section V offers a broader discussion of the findings and remaining questions, and Appendix A collects mathematical details. Three online appendices provide additional mathematical derivations, additional data analysis, and instructions to subjects.

I. Some Previous Work

Table 1 parametrizes the standard prisoner's dilemma payoff bimatrix (see Rapoport and Chamah 1965 for a related parametrization). With no loss of theoretical generality, the table normalizes the "cooperation" payoff at 10 and the "sucker" payoff at 0. Strategy B is strictly dominant, and so (B,B) is the unique Nash equilibrium, as long as the "temptation" payoff satisfies $x > 10$. The restriction $y < 10$ on the "punishment" payoff ensures that the Nash equilibrium is inefficient, and $x < 20$ ensures that the sucker-temptation profiles (A,B) and (B,A) also yield a lower payoff sum than the cooperation profile (A,A). Thus the dilemma: the unique equilibrium is inefficient.

A. Theory

The prisoner's dilemma was first posed by Melvin Dresher and Merrill Flood in 1950 and was given its famous framing story by Albert Tucker in that same year.¹ Since then, legions of theorists have sought ways to evade the dilemma and to support cooperation. John Nash apparently first pointed out (in private correspondence quoted in Flood 1952) that finite repetition doesn't help: all-defect—(B,B) in every period—remains the unique Nash equilibrium. Nash also pointed out that patient

¹Interestingly, the first published description of the game apparently was an experimental paper by Merrill M. Flood (1952), reporting an informal 100-period repeated prisoner's dilemma session with asymmetric temptation payoffs.

TABLE 1—GENERIC FORM OF PRISONER’S DILEMMA, WITH $20 > x > 10 > y > 0$

	A	B
A	(10, 10)	(0, x)
B	(x, 0)	(y, y)

pairs of players matched over an infinite sequence of stages can support cooperation, for instance by implementing grim trigger strategies. However, by the Folk Theorem (e.g., Fudenberg and Maskin 1986), they can just as easily support (as a Nash equilibrium of the repeated game) all-defect and a wide variety of other inefficient profile sequences.

What happens if the game is played over a continuous finite time interval, say $t \in [0, 1]$? Perhaps the most obvious approach is to specify a minimum reaction time τ and to formalize the game as a finitely repeated game with $1/\tau$ stages (rounding up to the nearest integer). The theoretical prediction again is that the dilemma persists, and only all-defect survives in Nash equilibrium.

Huberman and Glance (1993) show that cooperation evaporates in spatial versions of the repeated prisoner’s dilemma when players move asynchronously, in real time. According to the authors, the lesson is that coordination and cooperation can get an artificial boost when all players must move simultaneously at discrete time intervals. Clearly there is a need for more fully articulated models of games played in continuous time.

James Bergin and Bentley MacLeod (1993) develop one such model. They assume that actions cannot be reversed within time ϵ , look for ϵ -equilibria, and pass to the limit as ϵ goes to 0. For the prisoner’s dilemma, they obtain a Folk Theorem result: virtually any profile sequence that gives each player an average payoff of at least $y < 10$ can be supported as a Nash equilibrium (indeed, one that is renegotiation-proof).

Radner (1986) had previously studied ϵ -equilibria of the finitely repeated prisoner’s dilemma and (although he did not emphasize it) obtained an insight that we find very useful. Assume that players seek to maximize the average payoff over $T < \infty$ repetitions of the game in Table 1. Let C_k denote the strategy of playing Grim (i.e., choosing A until the other player plays B and choosing B thereafter) up through period k and playing B in the remaining periods $k + 1, \dots, T$. The usual unraveling argument notes that C_{k-1} is a best response to C_k for all $k = 1, \dots, T$. However, for $\epsilon > (x + y - 10)/T$, Radner’s equations (17, 19) show that C_{T-1} yields a payoff within ϵ of the best response payoff against C_k for every $k = 1, \dots, T$. The insight is that, when T is large, you lose considerably more than ϵ if your defection time n is much earlier than the other player’s defection time k , but by choosing near-maximal n , you never lose more than ϵ , no matter how large or small is k . Thus, waiting longer to defect is nearly dominant in Radner’s setup. This suggests that mutual cooperation can prevail.

Simon and Stinchcombe (1989) propose a general model of games played in continuous time. They consider discrete grids in the time interval $[0, 1]$ for games with finite numbers of players and actions. Given some technical conditions (e.g., the number of strategy switches remains uniformly bounded for each player), in

the limit as the grid interval approaches zero they obtain well-defined games in continuous time. Subgame perfection is automatic in these games, but backward induction does not work because the real numbers are not well ordered (Anderson forthcoming). For example, time $t = 1$ has no immediate predecessor: for any previous time, say $t = 1 - h$, there are an infinite number of later times that fall before $t = 1$, e.g., $t = 1 - h/17$. Consequently, some repeated game equilibria disappear in the continuous limit, while new equilibria can appear. Consistent with Radner's insight, Simon and Stinchcombe focus on Nash equilibria which survive iterated deletion of weakly dominated strategies. For games in continuous time similar to the prisoner's dilemma, they find a unique such equilibrium outcome: full cooperation at all times.

Thus existing theoretical literature offers three competing predictions for our continuous time experiment. The (naively extended) theory of finitely repeated games predicts that the all-defect profile (B, B) will predominate; Simon and Stinchcombe (and Radner) predict that the full cooperation profile (A, A) will predominate; and the extended Folk Theorem predicts virtually any profile sequence that gives each player at least the all-defect payoff y .

This complete information literature predicts no other role for payoff parameters x, y within their admissible range. The famous Gang of Four model introduced by Kreps et al. (1982) incorporates a touch of incomplete information and obtains the qualitative prediction that cooperation will decrease when either x or y increases. The same prediction arises from Quantal Response Equilibrium (e.g., McKelvey and Palfrey 1995), from the heuristics of Rapoport and Chammah (1965), and from most other models that include some sort of noise or imperfect information.

B. Experiments

Rapoport and Chammah (1965) conducted laboratory experiments with variants of the prisoner's dilemma iterated over 350 stages with fixed pairs of subjects, changing the x, y parameters randomly between 50 stage blocks. They find mutual cooperation rates above 60 percent in blocks with lowest x, y parameters, and cooperation rates under 50 percent in blocks with highest x, y . Rapoport, Guyer, and Gordon (1976) fix $x = 15$, $y = 5$ in our normalization and report that, over 300 stages, individual cooperation rates initially declined, then rose modestly and averaged about 55 percent overall. Unfortunately, this early work with the finitely repeated prisoner's dilemma did not provide subjects the opportunity to learn the logic of backward induction, because there was no stationary repetition of the repeated game.

More recent experiments—e.g., Selten and Stoecker (1986), Andreoni and Miller (1993), Hauk and Nagel (2001), and Bereby-Myer and Roth (2006)—feature stationary repetition of ten-stage repeated prisoner's dilemmas; that is, each subject plays a sequence of different ten-stage games against different opponents. These papers report that, after several repetitions of the repeated game, most subjects cooperate in early stages, but cooperation begins to unravel around the fifth stage and is rare after the eighth stage. Thus, even with ample opportunity to learn, the unravelling process seems at best incomplete in the laboratory data.

Potential explanations include the sequential equilibria of Kreps et al. (which were motivated by the laboratory results), or widespread altruism. Russell W. Cooper et al. (1996) find both of these explanations inadequate. In their experiment, cooperation rates decline fairly steadily over periods, not abruptly as in pure sequential equilibria, and remain positive in the last period. Moreover, contrary to the best calibrated mixed sequential equilibrium, there is an increasing, not decreasing, rate at which cooperation declines in later subperiods (their Figure 3). Incomplete unraveling remains a puzzle.

Other experiments study the “infinitely repeated” prisoner’s dilemma, in which there is an announced probability q that the matching ends after the current stage. Roth and Murningham (1978) produced mixed evidence for the theoretical prediction that cooperation is possible in such games. More recently, Dal Bo (2005) finds that individual cooperation rates respond sensitively to q , exceeding 50 percent in the most favorable case, while Aoyagi and Frechette (2009) observe individual cooperation rates as high as 85 percent under a very high q of 0.9.

Dal Bo and Frechette (2011) show that experience in these repeated games does not necessarily lead to greater cooperation and that variation in parameters similar to our x and y has a significant impact on cooperation. Individual cooperation rates average roughly 35 percent and rise to 76 percent with parameters more conducive to cooperation than any used in our own study. These authors conclude that the “shadow of the future” seems pivotal to cooperation.²

There are several ways to extrapolate these empirical results to continuous time. In 60-second periods, the shadow of the future shrinks steadily to zero. Will cooperation also decline steadily to zero? With more than 30 stationary repetitions per session and continuous time, our experiments provide unusually good learning opportunities. Will subjects learn to unravel cooperation more completely? Or will the absence of a well ordering, or asynchronous decisions, or other aspects of continuous time, twist the strategic behavior in a different direction? Answering such questions clearly requires new experiments.

II. Treatments and Experimental Design

We ran experiments using a new software package called ConG, for Continuous Games. Figure 1 shows the user interface. Each subject can freely switch between row actions A and B by clicking a radio button (or pressing an arrow key), causing the chosen row to be shaded. In our main treatment (Continuous time) the other player’s current choice is shown as a shaded column, and the intersection is doubly shaded. The computer response time to action switches is less than 50 milliseconds, giving players the experience of continuous action. The screen also shows the time series of actions (coded here as 1 for A and 0 for B) for the player and her counterpart in the upper right graph, while flow payoffs for each player are shown in the lower right graph. The top of the screen also shows the time remaining and the accumulated flow payoff.

²Indeed, under very favorable payoff parameters and a very strong shadow of the future, Dal Bo and Frechette eventually observe individual cooperation rates as high as 96 percent.

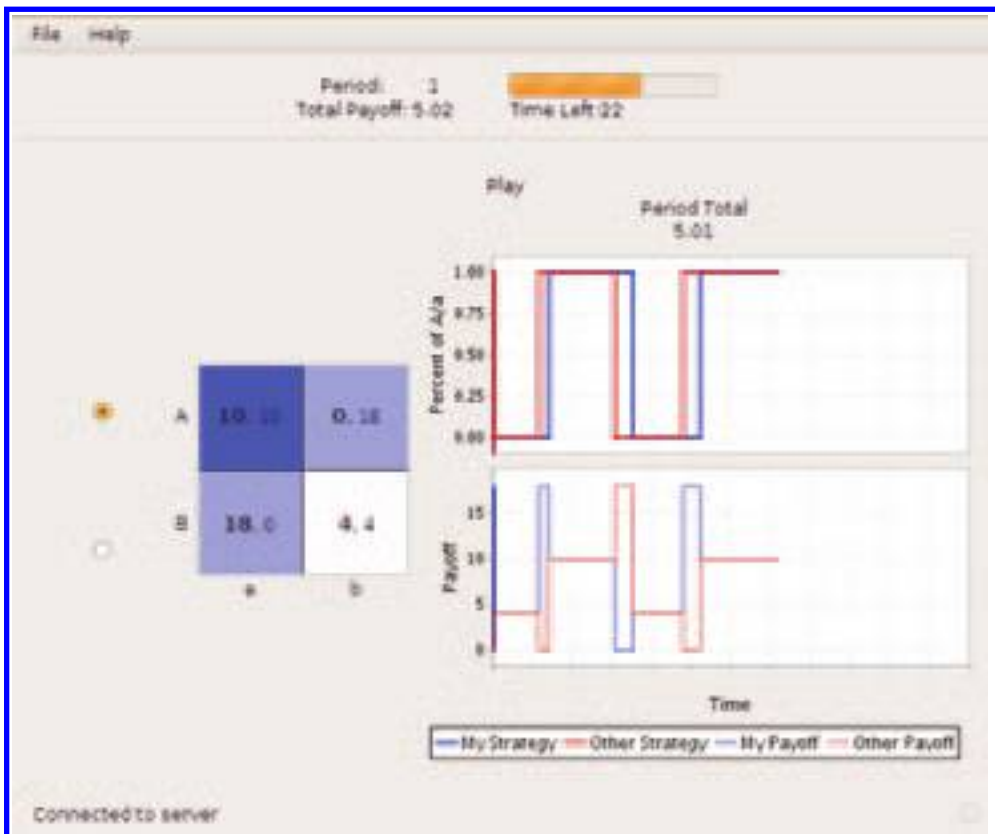


FIGURE 1. SCREENSHOT OF CONTINUOUS TIME DISPLAY

We study three treatments of time: Continuous, One-Shot, and Grid. In all treatments, each period lasts 60 seconds, during which subjects are allowed to change their actions at will. In Continuous time, subjects observe the unfolding history of actions and payoffs, and at the end of the period they earn the integral of the flow payoffs shown in the lower right hand graph of Figure 1.

In One-Shot time, subjects do not observe their counterpart's action until the period's end. They earn the lump sum payoffs for the action profile chosen at that point.

Grid time divides each 60-second period into n equal subperiods. Payoffs in each subperiod are determined only by the last action profile chosen in that subperiod. Only at the end of the subperiod does a player see her counterpart's choice, and that last profile becomes the initial profile of the next subperiod. Payoffs for the entire period are the average of the lump sum subperiod payoffs or, equivalently, the integral across subperiods of the piecewise constant flow payoffs. Thus One-Shot time is the same as Grid time with $n = 1$, and Continuous time is closely approximated by Grid time with $n > 300$.

Our other treatment variable, payoff parameters, examines four different configurations of (x,y) , one from each quadrant of the admissible domain $(10, 20) \times (0, 10)$.

They are Easy = (14, 4), Mix-a = (18, 4), Mix-b = (14, 8), and Hard = (18, 8). The names reflect the presumption that cooperation will be more difficult given either a larger temptation x or a larger punishment payoff y .³

In all treatments, subjects are randomly rematched with a new counterpart each period. At period's beginning, each of the four possible initial profiles is chosen independently with probability 0.25. Within period, profiles are automatically extended forward in time (to the next subperiod in Grid) until a player changes her action.

We first ran four sessions for Continuous and three parallel sessions for One-Shot.⁴ Ten subjects participated in each session (except for one Continuous session with only eight subjects), which consisted of 32 periods divided into eight blocks. Each of the four parameter sets appears once, in random order, in each block, and the sequences are matched across the two time treatments. Then we ran another four matched sessions, again using the same sequences, under the Grid treatment with $n = 8$ subperiods (hereafter called Grid-8 sessions). This treatment is comparable to the 10-stage repeated games featured in previous laboratory studies. We also ran three additional Grid sessions, to be described later, that varied n within session.

A key aspect of our design is that period lengths and potential payoffs are kept constant across Continuous, One-Shot, and Grid treatments. The only difference between these treatments is the frequency with which subjects can adjust their payoff-relevant choices.

Subjects in all sessions were randomly selected using online recruiting software at the University of California, Santa Cruz from our pool of volunteers, undergraduates from all major disciplines. They were all inexperienced, i.e., had never participated in a prisoner's dilemma experiment in our lab. On arrival, subjects received written instructions (available in online Appendix D) which also were read out loud. Sessions lasted on average 75 minutes, subjects were paid 5 cents per point each period, and average earnings were roughly US\$17.50 per subject.

III. Main Results

To provide an overview, we compile the fraction ρ_{ipk} of time spent in each profile ρ by player i and her counterpart $j(i, p, k)$ in period p of session k . Due to the symmetry of the game, the four action profiles reduce to three player-pair profiles:

- Mutual Cooperation ($\rho = c$): Profile (A, A);
- Mutual Defection ($\rho = d$): Profile (B, B);
- Sucker-Temptation ($\rho = s$): Profile (B, A) or (A, B).

For example, if player 2 spent equal time in each of the four action profiles in period 3 of session 4, the player-pair profile data would be $c_{234} = d_{234} = 0.25$ and $s_{234} = 0.50$. By definition, $\rho_{ipk} \in [0, 1]$ and $\sum_{\rho} \rho_{ipk} = c_{ipk} + d_{ipk} + s_{ipk} = 1$. Of course, $\rho_{ipk} = 0$ or 1 in the One-Shot treatment, and $\rho_{ipk} \in \{m/n : m = 0, 1, 2, \dots, n\}$ in the Grid- n treatment.

³ After normalizing, the payoff parameters used in the studies mentioned in the previous section mostly are in the neighborhood of our Easy parameters, and only a few are more challenging than our Mix parameters.

⁴ A coding error garbled several periods in a single One-Shot session. The data analysis to follow drops these periods, but all results are robust to, instead, using the entire dataset or dropping the entire offending session.

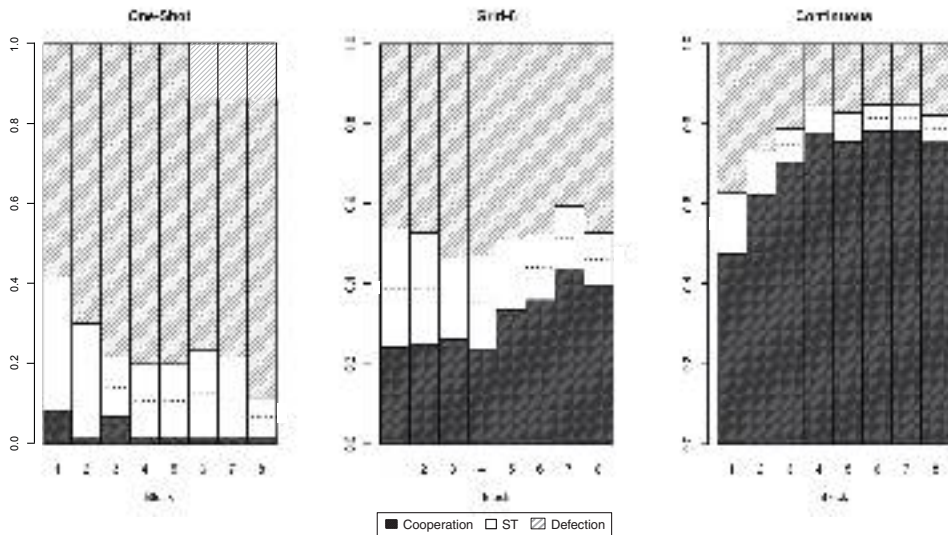


FIGURE 2. OUTCOMES OVER BLOCKS BY TREATMENT

Figure 2 shows mean rates of the three player-pair profiles $\rho = c, d, s$ over successive four-period blocks,

$$(1) \quad \rho_{bT} = \frac{\sum_{k \in T} \sum_{p \in b} \sum_i \rho_{ipk}}{\sum_{\rho} \sum_{k \in T} \sum_{p \in b} \sum_i \rho_{ipk}}$$

The randomized block design ensures that each block $b = 1, \dots, 8$ includes an equal sample of each of the four parameter sets for each time treatment $T = \text{One-Shot, Grid-8, Continuous}$.

Behavior seems fairly settled after block 3 (period 12). From this point onward, mean mutual cooperation rates are zero in One-Shot but approach 80 percent in Continuous. The mutual cooperation rate in Grid-8 is intermediate at about 30 percent. Mutual defection rates have the opposite pattern, since mean s rates are low in One-Shot and Grid-8 and are even lower in Continuous.

Most previous authors report *individual* cooperation rates $\kappa = c + 0.5s$, which can be seen in Figure 2 as horizontal dotted lines that bisect the ST bar. Our subsequent analysis will focus on settled behavior; unless otherwise noted the data will be drawn from blocks 4–8 (periods 13–32). None of the broad conclusions is altered by including the noisier data from blocks 1–3, but in some cases the statistical significance is lower.

Cumulative distribution functions (CDFs) reveal heterogeneous behavior. In Figure 3 we plot CDFs of mutual cooperation rates across periods and pairs for each treatment. In the Continuous treatment, the cooperation rate exceeds 80 percent for about two-thirds of the pairs, while in One-Shot cooperation virtually never happens. In Grid-8, a plurality (not quite a majority) of periods have a mutual cooperation rate of zero, but a substantial minority (about a third) have rates of 75 percent or more.

For none of these treatments is behavior clustered symmetrically around the mean. The median therefore provides a more reliable measure of central tendency, and it will be our focus for the remainder of the data analysis. Table 2 shows the median




FIGURE 3. CDFs OF COOPERATION RATES

TABLE 2—MEDIAN COOPERATION RATES (and bootstrapped standard errors)

Parameters	x	y	Continuous	Grid-8	One-Shot
Easy	14	4	0.931 (0.014)	0.750 (0.066)	0.000 (0.000)
Mix-a	18	4	0.890 (0.012)	0.500 (0.118)	0.000 (0.000)
Mix-b	14	8	0.905 (0.013)	0.000 (0.028)	0.000 (0.000)
Hard	18	8	0.811 (0.028)	0.000 (0.005)	0.000 (0.000)
All			0.893 (0.009)	0.250 (0.105)	0.000 (0.000)

cooperation rate in each treatment cell (again, for periods 13–32). Results are striking. For each parameter set, these rates are all zero in One-Shot, but in Continuous they range from 81 percent (in Hard) to over 93 percent (in Easy). The cooperation rates in Grid-8 are far more heterogeneous, ranging from zero (in Hard) to 75 percent (in Easy). Overall, as can be seen in the bottom row, there is a strong increase in cooperation as we move from One-Shot to Grid-8 to Continuous, and pairwise Mann-Whitney tests applied to by-subject median cooperation rates confirm this ordering at the 1 percent level.

To summarize,

RESULT 1: *Cooperation prevails in the Continuous treatment, is less than half as common in Grid-8, and is quite rare in One-Shot.*

TABLE 3—QUANTILE REGRESSION COEFFICIENT ESTIMATES
(and bootstrapped standard errors) FOR EQUATION (2)

Variable	Continuous	Grid-8
Intercept	0.919*** (0.018)	0.75*** (0.052)
X	-0.032 (0.021)	-0.25** (0.119)
Y	-0.013 (0.024)	-0.75*** (0.103)
$X \times Y$	-0.057 (0.043)	0.25* (0.149)

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

Table 2 also suggests that the impact of parameters differs across the time treatments. In Grid-8, cooperation never takes hold when $y = 8$ (Mix-b and Hard), but it is substantial when $y = 4$. In this last case, the x parameter also seems to have an impact. Parameters have no visible impact in One-Shot and a relatively small impact in Continuous, probably because cooperation rates are already so extreme in those treatments.

To follow up on these impressions, we ran quantile regressions of the form

$$(2) \quad c_{ij} = \beta_0 + \beta_x X_{ij} + \beta_y Y_{ij} + \beta_{xy} X_{ij} \times Y_{ij} + \epsilon_{ij},$$

where c_{ij} is subject i 's median rate of cooperation (over periods) under parameter set j ; X and Y are indicator variables taking a value of 1 when x and y take their high values and ϵ_{ij} is a normally distributed error term. Table 3 reports separate estimates for Continuous data and Grid-8 data; there is insufficient variation in the One-Shot data to estimate the model. The intercept estimates the cooperation rate in the Easy treatment. Increasing either x or y does not significantly change cooperation rates in Continuous, but their joint effect (found by adding β_x , β_y , and β_{xy}) in the Hard treatment is significant at the 5 percent level. In Grid-8 both parameters are highly significant, both statistically and economically.

RESULT 2: *Parameters have large negative effects on cooperation rates in the Grid-8 treatment but have little or no effect in Continuous and One-Shot.*

A. Behavior within Continuous Periods

What forces support the remarkably high rate of mutual cooperation in the Continuous treatment? Some clues can be gleaned from the trends within periods. The top panel of Figure 4 plots median⁵ rates of mutual cooperation $c_{0.5}(t)$ at each second for each parameterization. The median initial rate $c_{0.5}(0)$ is zero due to the

⁵Mean rates are similar but less extreme—they rise more gradually, reach a lower plateau and begin to decline a few seconds earlier—because $c(t)$ is bounded above at 100 percent and choices are dichotomous.




FIGURE 4

Notes: The top panel shows median rates of cooperation on a one-second grid in Continuous. The bottom panel shows cooperation rates at various quantiles of the distribution on the same one-second grid.

random assignments of initial actions $\rho_{ipk}(0)$, only a quarter of which are mutually cooperative. Strikingly, $c_{0.5}(t)$ rises to 100 percent by $t = 5$ seconds and remains there until only a few seconds remain. Then it falls rapidly, all the way to zero except in the Easy treatment.

The lower panel of Figure 4 aggregates across parameter sets but shows other quantiles $c_Q(t)$. The behavior at $Q = 0.85$, the 85th percentile, is mutual cooperation at each second, and $c_{0.75}(t)$ also is 1.0 except during the first two seconds. Lower quantiles indicate that cooperation ceases in the last few seconds for the majority of players. The graph of $c_{0.15}(t)$ shows that the cooperation ceases for more than 15 percent of the subjects when about 16 seconds remain (and doesn't begin for this fraction until 14 seconds have elapsed). The figure also shows that the cooperation

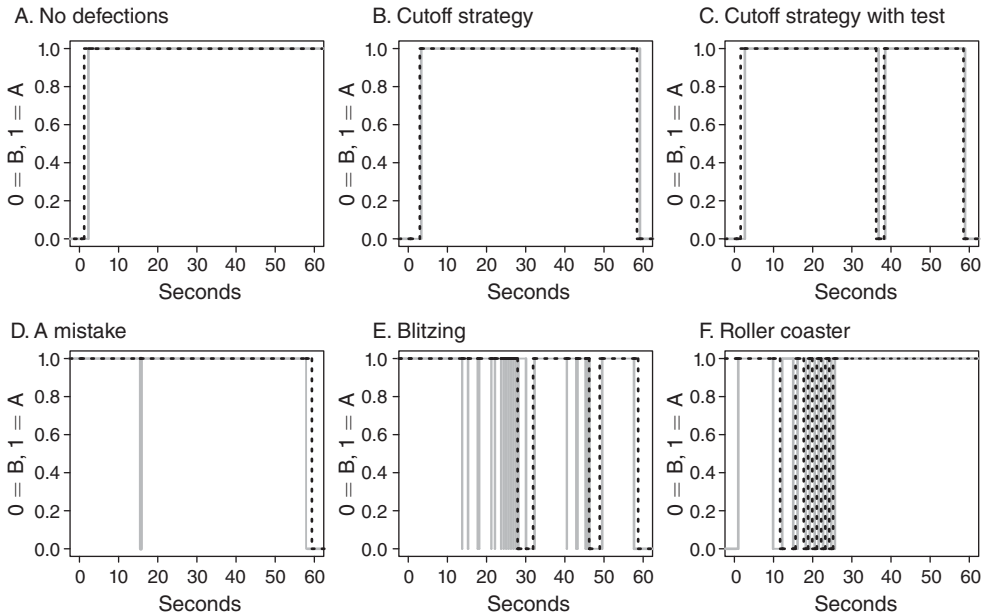


FIGURE 5. EXAMPLES OF WITHIN-PERIOD BEHAVIOR

level falls below 75 percent only when 5 seconds remain, and below 50 percent only when one second remains.

Figure 5 collects examples of the underlying individual behavior. Panel A shows one pair of players randomly initialized at the mutual defection profile $\rho = d$. At about $t = 2$ seconds, one of the players (marked in dotted heavy lines) switches to action A, putting the pair in Sucker-Temptation profile $\rho = s$. The other player (gray line) follows about a second later, and the pair remains in profile c the rest of the period. Player Gray earns 9.95 points, very close to the full cooperation payoff of 10, and player Dot is close behind at 9.7 points. Similar behavior by another pair is shown in panel B, except that Dot defects with about three seconds left in the period, and Gray follows within a second. Again, they earn just a bit under 10 points each (Gray 9.77, Dot 9.86).

Such behavior is quite typical. Ninety-five percent of player pairs initially assigned to profile $\rho = d$ or s sooner or later move to $\rho = c$, and the median time it takes to get there is only 2.89 seconds. As in panel (b), one of the players usually cuts off cooperation near the end of the period, and the other player quickly follows.

Table 4 shows that in the $20 \times 38 = 760$ player-periods observed in the settled final five blocks (20 periods) of the Continuous treatment, almost half of the observations include no defection from the mutual cooperation profile, and another 30 percent include only a single defection. About 9 percent defect twice, as in panel C of Figure 5. In this example, from an initial d profile player Dot switches within a couple of seconds to cooperate and player Gray quickly follows. As in the previous panels, profile c prevails for a time, but in this case, at about $t = 35$ seconds, Dot defects. Less than a second later, Gray follows suit. The resulting profile d doesn't last long. Dot soon switches back to A and Gray again follows quickly, restoring cooperation that lasts until almost the end of the period. With 2 or 3 seconds remaining (about the median, as shown in the last column of Table 4) Dot defects

TABLE 4—NUMBER OF BREAKS FROM MUTUAL COOPERATION, CORRESPONDING SHARE OF PLAYERS, AND CORRESPONDING MEDIAN TIME (*in seconds*) OF LAST MUTUAL COOPERATION

Number of breaks	Share	Cutoff time
0	0.479	60.0
1	0.305	57.6
2	0.089	57.9
3	0.057	57.3
4	0.022	57.6
5+	0.047	57.3

from mutual cooperation for the last time, and Gray again quickly follows. In this period, Dot earned 9.50 and, despite being twice suckered, Gray earned 9.60.

B. *Cutoff Strategies, Simple and Augmented*

Such behavior seems analogous to that observed in many repeated games settings, including Aoyagi and Frechette (2009), Dal Bo and Frechette (2011), and Engle-Warnick and Slonim (2006). To describe it in our setting, consider the following idealized strategy. Having achieved the mutual cooperation profile early in the period, a player using a *simple cutoff strategy* will unilaterally cut off cooperation only in the last few seconds but will match any prior defection as soon as possible. The median cutoff time in our data is with less than three seconds left, as shown in the last column of Table 4. The median response time to a defection is $\tau = 0.621$ seconds or about 1 percent of the period.⁶

Augmenting such simple cutoff strategies, players occasionally “test” the other player’s reciprocity reflexes (or attention), as Dot does at $t = 35$ seconds in panel (c). We refer to Dot’s subsequent return to cooperation as “repentance,” which is quickly accepted by Gray. Sometimes following a defection from cooperation, we see the other player offer “forgiveness” by switching briefly to action A; the defector usually follows quickly, restoring mutual cooperation. In an *augmented cutoff strategy*, a player will accept repentance, and may offer forgiveness, as long as his switch to A occurs before the chosen cutoff time and after sufficient time has passed that the other player does not profit from his defection.⁷

Almost 12 percent of the observations in Table 4 involve three or more defections, which suggests that not all observed behavior is consistent with these augmented cutoff strategies. Panel D of Figure 5 shows what seems to be a simple mistake: a player defected but returned almost instantly to the cooperation profile and remains there. In panel E, player Gray pulses briefly to defection at about $t = 13$ seconds, and returns before Dot reacts. This earns her a fraction of a penny, and she tries it again, and then many times again in the time interval (22, 28) seconds. At that point, Dot defects and Gray stops “blitzing.” A similar blitzing episode plays out more rapidly in the time interval (45, 49). A final pattern we call “rollercoastering”—both

⁶To be more precise, after filtering out the blitzes and roller coasters described below, the median duration of all profiles $\rho = s$ that follow $\rho = c$ is 0.621, and the mean is 0.91 seconds.

⁷Our working paper discusses this last point at greater length. A figure analogous to Figure 6 below, but including earnings in the subsequent profile, shows that in fact such defections are not profitable on average.

players blitz as in panel (f) and earn average flow payoffs $(10 + x + y)/4$, which is only 1 to 3 points short of the cooperative rate of 10 points.

Since very few subjects consistently exhibit any such behavior, we regard these deviations from cutoff strategies mainly as either brief errors or inexpensive escapes from boredom. Perhaps a stronger piece of evidence is the fact that following defection from mutual cooperation prior to the last 10 seconds, the defector repents in 82 percent of cases, and the other player offers forgiveness in another 16 percent. In only 1.7 percent of these cases do players simply stay in defection, and they eventually resume mutual cooperation in 85 percent of cases. Thus most deviations from cooperation before the last few seconds are quite transitory.

RESULT 3: *Subjects in the Continuous treatment tend to use simple or augmented cutoff strategies, defecting near the end of the period. Earlier deviations from mutual cooperation are typically short lived and are followed by a return to mutual cooperation.*

Observed behavior in the Grid-8 treatment can also be interpreted as arising from cutoff strategies. Conditional on reaching a cooperative profile, the median Grid-8 subject defects during a cooperative profile only once.

Are cutoff strategies an expensive luxury? In the Continuous treatment, defections from mutual cooperation are usually matched so quickly that, prior to the last few seconds, they do not seem to improve expected earnings. Likewise, unilateral moves to cooperation can be matched (or retracted) so quickly that they cost very little relative to staying in d . Figure 6 plots CDFs of the direct payoff impact of switching from mutual cooperation and from mutual defection, relative to the counterfactual of staying put. It shows that the costs of initiating cooperation and the benefits of initiating defection are nearly always below one cent and do not vary much across parameters. The vertical dashed lines in the figure show that, by contrast, the payoff effects in Grid-8 are much larger and differ substantially across parameters.

IV. Second Round Predictions and Experimental Results

The last section distilled from the data an intuitively appealing explanation for the high rates of mutual cooperation. In continuous time, the temptation to defect from $\rho = c$ nearly vanishes when a player expects her counterparty to reciprocate immediately. Likewise, at $\rho = d$, signaling a willingness to cooperate has negligible opportunity cost. Augmented cutoff strategies therefore are quite economical. They seem quite prevalent in our data and clearly are capable of supporting very high rates of cooperation. In our Grid-8 treatment, and in previous experiments with the finitely repeated prisoner's dilemma, the same forces are attenuated but still might support the observed moderate rates of cooperation and heterogeneity across parameter sets.

Intuitive appeal is a good start, but the explanation raises several questions. Why don't the cutoff strategies unravel? Could they constitute some sort of equilibrium? Does the explanation have any testable implications? Are there connections to existing theory?

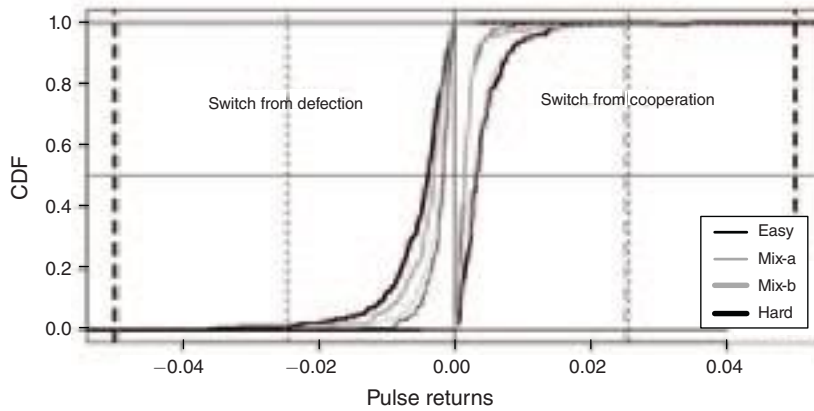


FIGURE 6

Notes: Empirical cumulative distributions of returns from switching from mutual defection and from mutual cooperation. Return is the flow payoff in dollars accumulated by the switcher until the next switch by either player, less the flow payoff that would have been accumulated in the original profile over the same time interval.

Our answers to these questions are inspired by Radner (1986), who shows that highly cooperative outcomes can be supported as part of an equilibrium if agents are willing to deviate slightly from best responses. Appendix A adapts Radner's model to a continuous time interval normalized on $[0,1]$ with small reaction lags $\tau \geq 0$. Like Radner (1986), our analysis focuses on a very particular subset of ϵ -equilibria that can generate a high level of cooperation. Of course, as with Radner's model, many other ϵ -equilibria can be constructed for this game, so ours is not the definitive analysis.

To achieve our focus, we impose a number of strong simplifying assumptions that we do not fully justify but that seem consistent with the data. First, we follow Radner in assuming that agents cooperate initially and in restricting the strategy space to simple cutoff strategies. Though restrictive, this assumption broadly matches our empirical findings—recall that subjects usually achieve cooperation quickly and, in nearly 80 percent of cases, they defect no more than once. Second, we assume that all players share the same reaction lag $\tau \geq 0$.

Third, because subjects frequently cooperate until the end, we assume that subjects are insensitive enough to payoff shortfalls that we don't rule out full-time conditional cooperation, i.e., waiting until the end before unilaterally cutting off mutual cooperation.⁸ Radner's (1986) key insight is that full-time conditional cooperation is almost a dominant strategy in a long finite horizon game. To see that this insight extends to a continuous time setting with quick reaction times, consider the following thought experiment.

Suppose that your opponent plays a simple cutoff strategy $K(u)$ with cutoff time $u \in (\tau, 1)$. Your best response, of course, is to play $K(s)$ with $s = u - \tau$, but suppose that you instead cut off cooperation too early, at time $s = u - z > 0$, with $z \geq \tau$. Relative to the best response, you lose $(10 - y)(z - \tau)$ because over a time interval of length $z - \tau$ you get the punishment payoff y instead of the cooperation

⁸ A notable consequence of this assumption is that the required degree of insensitivity, ϵ , depends on the size of the reaction lag, τ .

payoff 10. This loss represents a substantial fraction of potential earnings when z is substantially greater than τ .

On the other hand, suppose that you plan to cut off cooperation too late, say at time $s = u + z \leq 1$ for some $z \geq 0$. Then relative to the best response you earn the sucker payoff 0 instead of the defection payoff y until you react to your opponent's defection, and you also forgo the temptation $x - 10$ available before your opponent reacts to your defection. Hence your loss is $(x + y - 10)\tau$. Crucially, this loss is independent of z ; it depends only on the reaction speed and the payoff parameters. It follows that, no matter which cutoff time u your opponent selects, it is an ϵ -best response to set $s = 1$ as your cutoff time, for any $\epsilon \geq (x + y - 10)\tau$. Thus $K(1)$, full-time conditional cooperation, is nearly a dominant strategy when τ is small.

Section A1 of Appendix A extends this sort of argument to the case where a player is uncertain of his opponent's cutoff time u but knows the distribution from which it is drawn. Proposition 1 of the Appendix shows that $K(1)$ becomes an undominated strategy in an appropriate limit as $\tau \rightarrow 0$. Thus Radner's insight allows us to recover something reminiscent of Simon and Stinchcombe's result.

Section A2 of Appendix A looks for ϵ -equilibria when τ is fixed at a small but positive value. Such equilibria take the form of a mixture of cutoff times,⁹ such that each cutoff time potentially used is nearly a best response to the mixture. The set of such equilibria seems to be quite large, so we focus on strategies that are not dominated by the simple and focal strategy, $K(1)$. This strategy of full-time conditional cooperation is the analog in our setting of the famous grim trigger strategy, and mixes that remain after deleting strategies dominated by $K(1)$ are called ND distributions.

Proposition 2 of Appendix A shows that any ND distribution constitutes an ϵ -equilibrium for any $\epsilon \geq (x + y - 10)\tau$. Appendix A notes that it is realistic empirically and sensible theoretically to assume that F is not negatively skewed, and to assume that (except perhaps for mass points at $s = 1$ and $s = 0$) it has a density f . For this case it shows that all strategies employed in any ND ϵ -equilibrium involve cutoffs after time

$$(3) \quad s_L = 1 - \frac{2x}{10 - y} \tau.$$

Under the stronger assumption that cutoff times are uniformly distributed on $[s_L, 1]$ with $s_L > 0$, Appendix A derives the prediction that the median overall fraction of cooperation will be

$$(4) \quad c_{0.5} = 1 - \frac{\sqrt{2}x}{10 - y} \tau.$$

Online Appendix B shows that the predicted level of cooperation is arguably zero when the expression for s_L is negative in (3) and is surely zero when the expression for $c_{0.5}$ is

⁹Due to the strategic symmetry of the prisoner's dilemma, one looks for a single mixture, not a pair of mixtures. Given asynchronous choices, it seems infeasible for players to coordinate on precise interior clock times, so Proposition 2 assumes that mixes are nondegenerate over $t \in (0, 10)$.

TABLE 5—RATES OF COOPERATION PREDICTED BY EQUATION (4), OBSERVED MEDIAN RATES AND OBSERVED FINAL TIMES OF MUTUAL COOPERATION (with bootstrapped standard errors).

Parameters	x	y	Predicted	Cooperation median rate	Final time
<i>Panel A. Continuous</i>					
Easy	14	4	0.967	0.931 (0.014)	0.994 (0.009)
Mix-a	18	4	0.958	0.890 (0.012)	0.973 (0.007)
Mix-b	14	8	0.901	0.905 (0.013)	0.974 (0.006)
Hard	18	8	0.873	0.811 (0.028)	0.959 (0.006)
<i>Panel B. Grid-8</i>					
Easy	14	4	0.588	0.750 (0.066)	0.750 (0.054)
Mix-a	18	4	0.470	0.500 (0.118)	0.625 (0.092)
Mix-b	14	8	0.000	0.000 (0.028)	0.000 (0.032)
Hard	18	8	0.000	0.000 (0.005)	0.000 (0.006)

negative in (4). We should reiterate that these equations come from a particular class of ϵ -equilibria, and there exist other sorts of equilibria with different predictions. However, as we show below, this class describes our data remarkably well.

In our Continuous treatment, τ depends on subjects' endogenous reaction time, which we estimate to be 0.01 minutes. With τ that small, equations (3–4) predict that cutoff times occur very near the end of the period for all (x, y) parameters that we used. The intuition is that the temptation to defect and the risks from suffering defection are quite small in this case. By contrast, in the Grid-8 treatment, τ is exogenous: subjects are forced to wait $\tau = 0.125$ of the period to react to a unilateral defection by a counterpart. The equations then predict considerably earlier cutoff times that vary substantially with the (x, y) parameters. Of course, with $\tau = 1.0$ in the One-Shot treatment, the equations give negative values, and the prediction is no cooperation.

These predictions match well the patterns observed in the data. Table 5 calculates (4) under each parameter set using $\tau = 0.01$ and $\tau = 1/8$ for Continuous and Grid-8 respectively. It also reproduces median final cooperation times (a literal interpretation of the prediction) and median overall cooperation rates for Continuous and Grid-8. The predictions capture the comparative statics reported in Result 2—the impact of parameters (especially of y) is large in Grid-8 but is modest in Continuous. Even the point predictions are quite close to the data, except that the Grid-8/Easy observations are about 16 percentage points (a bit more than one subperiod) above the prediction.

RESULT 4: *Parameter effects in Continuous and Grid-8 are rather well explained by a class of epsilon equilibria in cutoff strategies.*




FIGURE 7

Notes: Grid- n predictions and data. Predictions are from equation (4) with $x = 18$; $y = 4$ and $\tau = 1/n$ for $n = 2, 4, 8, 15, 30, 60$. Data are median times of final mutual cooperation and median rates of mutual cooperation. Horizontal lines show Mix-a data for One-Shot and Continuous sessions.

A. Grid- n Sessions and Results

A class of ϵ – equilibria in cutoff strategies governed by reaction lags $\tau > 0$ accounts nicely for the data so far. Of course, the true test of any explanation lies in its excess predictive power—in verified implications beyond the facts it was constructed to explain. The key prediction of the model is that cooperation rates rise from zero (observed in One-Shot) to about 90 percent (observed in Continuous) as the forced reaction lag, τ , shrinks to its minimal feasible value for humans.

To better test this prediction, we conducted additional laboratory sessions that exogenously controlled τ at multiple levels. In a Grid- n treatment, the number n of subperiods exogenously imposes the minimum reaction time $\tau = 1/n$. Equation (4) predicts a monotonic (indeed linear) decrease in cooperation as τ rises from zero.

Figure 7 plots the theoretical predictions for Mix-a parameters and $\tau = 1/n$ as solid black lines connecting the dots for $n = 60, 30, 15, 8, 4, 2$ subperiods. For reference, the horizontal gray lines plot median rates of cooperation observed in Mix-a Continuous and One-Shot ($\tau = 1$) periods. The prediction is that as the grid gets coarser, cutoff times fall from nearly continuous levels at $n = 60$ to One-Shot levels at $n \leq 4$.

We ran three sessions of Grid- n , each lasting 36 periods and using only the Mix-a parameters. In each of three 12-period blocks, we ran each n twice in consecutive periods and varied n in the sequence Incr = (2, 4, 8, 15, 30, 60) or Decr = (60, 30, 15, 8, 4, 2) or Random. (In two sessions the blocks were sequenced

Incr-Decr-Random, and in the other session the blocks were sequenced Decr-Incr-Random.) Note that the within session variation of n allows us to observe each subject's behavior at each value of τ , generating a particularly stringent test. To focus on settled behavior, we once again examine data after the first 12 periods (after subjects have experienced each n twice); the results are similar (though a bit noisier) if we include all data.

The median subject in the Grid- n sessions employed a cutoff strategy, departing from the c profile only once per period, just as in Continuous and Grid-8. As predicted, the observed timing of cutoffs depended negatively on n . Figure 7 plots the median final cooperation time and the median rate of cooperation as a function of $\tau = 1/n$. When periods are divided into 2 or 4 subperiods ($\tau = 0.5$ or 0.25) cooperation never gets off the ground. However, as the grid gets finer and reaction lags become smaller, cooperation rates start to rise towards Continuous levels at $n = 60$ ($\tau = 0.016$). Most strikingly, the observed median rates of cooperation and cutoff times tightly bracket the point predictions of the model.

RESULT 5: As the grid becomes finer in Grid- n sessions, the threshold times and cooperation rates approach those observed in Continuous. Moreover, both median cutoff times and cooperation rates fall nearly linearly with τ and closely track the point predictions given by equation (4).

V. Discussion

Our principal findings can be summarized briefly. First and foremost, in the Continuous time treatment, we found remarkably high levels of mutual cooperation in all four parameterizations of the prisoner's dilemma. Even with "hard" parameters (maximal temptation and minimal efficiency loss), the all-cooperate profile was played 81 percent of the time by the median pair of subjects in later periods. The other parameterizations led to median mutual cooperation rates of 89 to over 93 percent. By contrast, in the Grid-8 control treatment, with rapid repeat pairings over eight subperiods, defection was more prevalent than cooperation, and cooperation was rare in the One-Shot control treatment.

Second, the parameterization had a considerably stronger impact (in the predicted direction) in the Grid-8 treatment than in Continuous time. In the One-Shot treatment, parameters had negligible impact because cooperation was always rare.

Third, within the 60 second Continuous periods, median rates of cooperation quickly reached 100 percent and remained there until the last few seconds of the period, when they dropped off abruptly.

Inspired by a strand of existing theoretical literature, we postulated a particular class of epsilon equilibria and derived formulas predicting how cooperation rates respond to adjustment lags and to payoff parameters. These predictions accounted well for the Continuous, Grid-8, and (trivially) One-Shot data. They also nicely explained a set of second-round data from Grid- n sessions, which varied the number of subperiods from 2 to 60. Thus the formulas correctly predict defection in one-shot games, cooperation in continuous time and intermediate results on the path between the two.

The underlying intuition is simple. When your opponent can react very quickly, defecting from mutual cooperation is likely to earn you the temptation payoff only briefly and may cost you the cooperation payoff for the rest of the period. Likewise, briefly switching from mutual defection is cheap for you, and may catalyze a sustained move to a higher payoff profile. Hence rapid reactions tend to stabilize mutual cooperation and destabilize mutual defection, at least until late in the period.

Conventional wisdom is that cooperation is susceptible to unraveling when the time horizon is finite. Sufficiently late in the period, your incentive is to defect. Your opponent's incentive is to defect before you do, and yours is to defect before she does, so backward induction might seem to unravel cooperation. However, experimentalists since Selten and Stoecker (1986) have found that, even given good learning opportunities, unraveling typically doesn't actually go very far.¹⁰

Our work shows that the unraveling argument loses its force when players can react quickly. The faster she can react, the less incentive you have to preempt your opponent, and the earlier you defect, the more you stand to lose from preempting her. Extending the ideas of Radner (1986), and Simon and Stinchcombe (1989), we took some first steps towards formalizing this argument in terms of epsilon equilibrium. Unraveling is quite limited when players are willing to sacrifice a small part of their potential payoff and they can react sufficiently rapidly. The faster they can react, the smaller the potential sacrifice and the greater the level of cooperation.

Our results set the stage for new theoretical advances. Although the epsilon equilibrium analysis in Appendix A organizes and explains our empirical results, it rests on several simplifying assumptions that are not fully justified. Additional insights may be gained by a more definitive analysis, or by considering alternative approaches. More broadly, as noted in Section I, strategic interaction in continuous time may be affected by asynchronicity and by the fact that the real numbers are not well ordered. These features of continuous time play only a minor part in Appendix A, but new theoretical analysis may find more substantive roles for them.

Much empirical work also remains. Future laboratory studies could test robustness of our predictions to different payoff parameters, to longer or shorter periods or more periods, and to variations on near continuous time, e.g., alternating moves, or perceptible lags in implementing action switches, or temporary action lock-ins. More generally, further studies might examine whether continuous time can ever reduce efficiency,¹¹ and seek additional practical insights into the forces that encourage or discourage efficient cooperation.

¹⁰Nash anticipated this result, as well as our augmented cutoff strategies. According to Flood (1952, footnote on p. 24), Nash wrote of the 100-period iterated prisoner's dilemma "... one should expect an approximation to [grim trigger] ... with a little flurry of aggressiveness at the end and perhaps a few sallies, to test the opponent's mettle during the game."

¹¹See Anderson, Jenny. 2009. "US Proposes Ban on 'Flash' Trading on Wall Street." *New York Times*, September 17 (<http://www.nytimes.com/2009/09/18/business/18regulate.html>) for a possible practical example. In terms of matrix games, consider $x > 20$ in Table 1. Unlike in a true prisoner's dilemma game, the sucker-temptation profile is now efficient. In discrete time, efficiency might be achieved by alternating the two ST cells, and such coordination might be more difficult in continuous time.

APPENDIX A

This Appendix collects mathematical details. The goal is not to lay broad theoretical foundations for strategic interaction in continuous time given reaction lags, but rather to derive equations (3–4) of the text and to tighten connections to the theoretical work of Simon and Stinchcombe (1989) and Radner (1986).

A.1 Near Dominance

Consider continuous play of the prisoner's dilemma game with flow payoff shown in Table 1. Recall the restriction $20 > x > 10 > y > 0$ on the temptation and punishment parameters. Normalize the time interval to $[0, 1]$. In terms of the experiment, this corresponds to a single period measured in minutes, where the median reaction time was approximately $\tau = 0.01$ minutes.

Given a known reaction time $\tau > 0$, a *simple cutoff strategy*, denoted $K(s)$, specifies a time $s \in [0, 1]$ with unconditional defection (choosing action B) at all later times $t \geq s \in [0, 1]$ and conditional cooperation at all earlier times $t < s$. *Conditional cooperation* means that the player chooses action A until her opponent first chooses B, after which, with lag τ , she switches to B and remains there until time runs out. Of course, $K(0)$ represents unconditional and immediate defection, and $K(1)$ represents full-time conditional cooperation.

Let the space of pure strategies consist of all simple cutoff strategies with a given reaction time, $\mathcal{S}^\tau = \{K(s) : s \in [0, 1]\}$. A mixed strategy is represented by a cumulative distribution function $F(s)$ for the cutoff times $s \in [0, 1]$. The analysis below sometimes assumes that F has a well-defined density $f: [0, 1] \rightarrow [0, \infty]$. A justification is asynchronicity: it is infeasible for players to coordinate on specific clock times $t \in (0, 1)$; of course, they can coordinate on the endpoints $t = 0, 1$, so mass points may appear there.

The first task is to compute the expected payoff of an arbitrary cutoff time relative to full-time conditional cooperation. Let u represent the opponent's unknown cutoff time. A player choosing cutoff strategy $K(s)$ obtains payoff $[10u + 0\tau + y(1 - u - \tau)]$ if $u < s - \tau$, or $[10s + x\tau + y(1 - s - \tau)]$ if $u > s + \tau$. For $u \in [s - \tau, s]$ her payoff is $10u + 0(s - u) + y(1 - s)$, and for $u \in [s, s + \tau]$ it is $10s + x(u - s) + y(1 - u)$.

Given that potential opponents' strategies are drawn from mixture F with density f , the expected payoff for $K(s)$ with $s \leq 1 - \tau$ is

$$\begin{aligned}
 (5) \quad \pi(s|F) &= \int_0^{s-\tau} [10u + y(1 - u - \tau)] f(u) \, du \\
 &\quad + \int_{s-\tau}^s [10u + y(1 - s)] f(u) \, du \\
 &\quad + \int_s^{s+\tau} [10s + x(u - s) + y(1 - u)] f(u) \, du
 \end{aligned}$$

$$+ \int_{s+\tau}^1 [10s + x\tau + y(1 - s - \tau)] f(u) du.$$

The payoff for full-time conditional cooperation, $K(1)$, is the same over the region $u < s$ but differs over the region $s \leq u \leq 1$. Here, instead of the integrands in (5), the player gets $10u + y(1 - u - \tau)$. Thus the difference in expected payoff is

$$\begin{aligned} (6) \quad \pi(s|F) - \pi(1|F) &= \int_s^{s+\tau} [10(s - u) + x(u - s) + y\tau] f(u) du \\ &\quad + \int_{s+\tau}^1 [10(s - u) + x\tau - y(s - u)] f(u) du \\ &= -(10 - y) \int_s^1 (u - s) f(u) du + x\tau[1 - F(s)] \\ &\quad + (x - y) \int_s^{s+\tau} (u - s - \tau) f(u) du. \end{aligned}$$

Inspection of the last line of (6) suggests that it will be zero or negative when τ is sufficiently small. That is, full-time conditional cooperation seems to do better than other cutoff strategies in the limit, consistent with the results of Simon and Stinchcombe. More precisely,

PROPOSITION 1: *Fix $s \in [0, 1)$ and choose an arbitrary cumulative distribution function $F: [0, 1] \rightarrow [0, 1]$. Then $\lim_{\tau \searrow 0} [\pi(1|F, \tau) - \pi(s|F, \tau)] \geq 0$.*

PROOF:

To cover cases in which F has no density, rewrite (6) and preceding equations as Stieltjes integrals with $dF(u)$ everywhere replacing $f(u)du$. In the last line of (rewritten) equation (6), the first term is clearly nonpositive and independent of τ . The second term clearly vanishes as $\tau \searrow 0$. The integrand $(u - s - \tau)$ in the third term converges uniformly to zero as $\tau \searrow 0$, so this final term also vanishes in the limit. Hence, the entire expression reduces in the limit to the nonpositive first term. The conclusion follows immediately. QED

A.2 Equilibrium Predictions

The next task is to characterize a useful set of equilibrium distributions of cutoff times for fixed positive reaction lags τ . Recall that our focus is on ϵ -equilibrium, which allows for small shortfalls in expected payoffs relative to a best response. To formalize the idea, recall that $\text{Supp}F$, the *support of distribution F* , is the smallest closed set containing all points of increase of F . When F has density f , the support is the closure of $\{s: f(s) > 0\}$.

Given some $\epsilon \geq 0$, we say that a distribution F is an ϵ -equilibrium if, for every $s \in \text{Supp}F$, we have $\pi(s|F) \geq \sup_{x \in [0,1]} \pi(x|F) - \epsilon$. That is, in ϵ -equilibrium, every cutoff actually used is nearly a best response to the overall distribution.

We shall not try to characterize the full set of ϵ -equilibria but will focus instead on those that are (weakly) better responses than full-time conditional cooperation. Motivated by the discussion at the end of the previous section, we say that a distribution F representing a mixture of cutoff strategies is *nearly dominant* (ND) if it has a well-defined density over $(0, 1)$ and equation (6) is nonnegative for all $s \in \text{Supp}F$.

The following proposition justifies our focus. It says that ND distributions are in fact ϵ -equilibria for ϵ commensurate with the reaction time τ .

PROPOSITION 2: *Let $F: [0, 1] \rightarrow [0, 1]$ represent a ND mixture over the space \mathcal{S}^τ of pure cutoff strategies, and let s_L be the smallest solution to $\pi(s_L|F) = \pi(1|F)$. Then*

- (a) F is an ϵ -equilibrium for any $\epsilon \geq \hat{\epsilon} = (x + y - 10)\tau$;
- (b) $\text{Supp}F \subset [s_L, 1]$; and
- (c) $1 - s_L = A\tau + O(\tau^2)$, where $A > 0$ is given in a formula below involving the skewness of F and the payoff parameters x and y .

Recall that $O(x^n)$ is conventional notation for any term that is negligible of order n , i.e., after being divided by x^n the term remains bounded as $x \rightarrow 0$.

PROOF:

The discussion in the text preceding equation (3) establishes that $\hat{\epsilon} = (x + y - 10)\tau$ is a *uniform* upper bound on the payoff shortfall of $K(1)$ relative to any simple cutoff strategy, including the ex post best response—that is, the best response to the opponent’s realized cutoff time u earns at most $\hat{\epsilon}$ more than $K(1)$. It follows that $\hat{\epsilon}$ is an upper bound on the *expected* payoff shortfall of $K(1)$ relative to the ex ante best response. A trivial consequence is that it remains an upper bound for the shortfall of any strategy with expected payoff greater than $K(1)$. Part (a) follows immediately.

Part (b) is immediate from the definition of s_L and the evident continuity of (6) in s for $s \in (0, 1)$. If $s < s_L$, then the “smallest solution” property ensures that $\pi(s|F) < \pi(1|F)$ and so $s \notin \text{Supp}F$.

For part (c), note that in the last term in the last line of (6), the integrand expression $(u - s - \tau)$ is negative and has absolute value less than τ over the relevant interval, which has length τ . Thus this last term is $O(\tau^2)$ and negative. For notational simplicity, assume for the moment that this term is 0.

Part (b) ensures that $F(s) = 0$ for $s \leq s_L$, so over this range we can rewrite the payoff advantage (6) as

$$(7) \quad \pi(s|F) - \pi(1|F) = -(10 - y)\psi(s|F) + x\tau,$$




FIGURE 8. PLOT OF $\psi(s)$ AND THE DETERMINATION OF s_L

where $\psi(s|F) = \int_s^1 (u - s)f(u) du \in (0, 1 - s)$ has value \bar{u} at $s = 0$ and is a decreasing function with derivative $F(s) - 1$. The payoff advantage is clearly negative for such s as long as

$$\psi(s) > \frac{x\tau}{10 - y},$$

i.e., for $s < \psi^{-1}(x\tau/(10 - y))$. For $s \in [0, s_L)$, the function ψ has slope -1 and hence is simply $\psi(s) = \psi(0) - s$. As illustrated in Figure 8, we can now characterize the lower support point s_L by the equation $\pi(s_L) - \pi(1) = 0$ or, using (7),

$$(8) \quad \frac{x\tau}{10 - y} = \psi(s_L) = \psi(0) - s_L = \bar{u} - s_L.$$

Of course, the shape of the distribution F determines the position of \bar{u} within $[s_L, 1]$. We have $\bar{u} = \alpha s_L + (1 - \alpha)1$, where $\alpha \in (0, 0.5)$ represents an upward skew and $\alpha \in (0.5, 1)$ represents a downward skew. Substituting $\bar{u} - s_L = (1 - \alpha)(1 - s_L)$ into (8) and solving for s_L , we obtain

$$(9) \quad s_L = 1 - \frac{x\tau}{(1 - \alpha)(10 - y)}.$$

Repeating the analysis while keeping track of the $O(\tau^2)$ term in (6) yields messier expressions, but ultimately one obtains an analog of equation (9) with a (different) $O(\tau^2)$ term appended. The remaining technicality is that the hypothesis allows positive probability mass at the points, 0 and 1; it is routine to check that the argument still goes through in this case, using Stieltjes integrals where appropriate. Part (c) then follows, with constant $A = x/[(1 - \alpha)(10 - y)]$. QED

A few remarks are in order before proceeding further. In part (a), the ex post best response to the realized cutoff u is of course more profitable than the ex ante best response, which is constrained to be the same for all realized u . Hence, the expected shortfall for $K(1)$ is less than $\hat{\epsilon}$, and often considerably less. The upshot is that a ND distribution is an ϵ -equilibrium for a range of $\hat{\epsilon}$'s smaller than $\epsilon = (x + y - 10)\tau$, often considerably smaller.

In part (c), equation (9) reduces to (3) in the unskewed case $\alpha = 0.5$. As asserted in the text, this case represents a lower bound for $\alpha \in (0, 0.5)$, the upward skewed case.

We regard downward skew as less relevant, theoretically as well as empirically. Theoretically, the Radner argument shows that cutoffs below the midpoint are less robust to deviations than those above the midpoint. Empirically, we find that observed cutoffs have considerable upward skew. Nevertheless, the downward skew case is pedagogically useful. For any fixed positive values of τ, x, y , one can find a value of α sufficiently close to 1 so that the expression for s_L is zero (or negative). The interpretation is unraveling: when the distribution is sufficiently downward skewed, it pays to cut off cooperation earlier than the modal time, and cutoff times unravel all the way down to zero, as in the traditional analysis.

The next task is to predict the median fraction $c_{0.5}$ of time in mutual cooperation. Given the $[0, 1]$ time normalization, and the assumption of simple cutoff strategies initially in mutual cooperation, the fraction of cooperation coincides with the time of first defection. Hence it is the minimum y of two independent draws from the ND distribution F . Classic work on order statistics, e.g., Robert V. Hogg and Allen Craig (1970), shows that this minimum has density $g(y) = 2(1 - F(y))f(y)$. The median value $m = c_{0.5}$ of y therefore is the root of the equation $0.5 = \int_0^m g(y) dy$. Inserting the expression for g and dividing by 2, the equation becomes

$$0.25 = \int_0^m f(y) dy - \int_0^m F(y)f(y) dy = F(m) - 0.5 [F(m)]^2.$$

Multiplying through by 4, we see that $z = F(m)$ satisfies the quadratic equation $2z^2 - 4z + 1 = 0$, whose relevant root is $z = (4 - \sqrt{16 - 8})/4 = 1 - 1/\sqrt{2}$. Hence the median time of first defection, and thus the predicted median mutual cooperation rate, is

$$(10) \quad c_{0.5} = m = F^{-1}(z), \text{ where } 1 - z = \frac{1}{\sqrt{2}}.$$

When F is the uniform distribution on the interval $[s_L, 1]$ for $s_L \in (0, 1)$, then m can be written out explicitly. In this case, $z = F(m) = (m - s_L)/(1 - s_L)$. Multiplying by the denominator $1 - s_L$ and recalling from (3) that s_L takes the form $1 - A\tau$, we obtain

$$m = z(1 - s_L) + s_L = zA\tau + 1 - A\tau = 1 - (1 - z)A\tau = 1 - \frac{A\tau}{\sqrt{2}}.$$

Inserting $A = 2x/(10 - y)$ we obtain equation (4).

Thus, as a function of reaction time τ , the fraction of time in mutual cooperation is predicted to be linearly decreasing, with slope $-A/\sqrt{2}$, with A the given increasing function of the parameters x, y . For example, for the Mix-a parameters $x = 18, y = 4$, the slope is predicted to be $-36/6\sqrt{2} \approx -4.24$.

A.3 Gaps in the Theory

Online Appendix B completes two remaining tasks. It shows that two different ways of obtaining s_L in discrete time closely approximate each other and explains why negative values of s_L in (3) predict zero cooperation.

Important theoretical tasks remain unfinished. The current theoretical analysis does not fully justify the restricted strategy space and the focus on ND distributions. We conjecture that (in some sense that needs to be made precise) augmented cutoff strategies weakly dominate other feasible strategies when τ is negligible. Intuitively, the augmentation “testing” looks for exploitable lags in the opponent’s strategy, and the augmentations “repentance” and “forgiveness” render the opponent’s defection from $\rho = c$ unprofitable by requiring sufficient lags before returning to action A. Delete the other strategies in the first round, and restrict the strategy space to augmented cutoff strategies. Then we further conjecture that simple cutoff strategies become weakly dominant when τ is negligible. Intuitively, both testing and blitzing lose money against a simple cutoff strategy, and forgiveness is then redundant. If some such conjectures are correct, then the restricted strategy space would be justified by iterated dominance, roughly analogous to arguments used by Simon and Stinchcombe (1989).

Our justification for ND distributions in the space of cutoff strategies was ultimately empirical—they seem roughly consistent with observed behavior, and they suggested themselves to us due to the prominence of $K(1)$. It might be worth looking for other prominent epsilon equilibria in the space of cutoff strategies, and to see whether they could lead to predictions that differ substantially from ours. We hope that theoretically minded researchers will investigate these matters.

REFERENCES

- Anderson, Robert M.** Forthcoming. “Quick Response Equilibria.” *Advances in Theoretical Economics*.
- Andreoni, James A., and John H. Miller.** 1993. “Rational Cooperation in the Finitely Repeated Prisoner’s Dilemma: Experimental Evidence.” *Economic Journal* 103(418): 570–85.
- Aoyagi, Masaki, and Guillaume Frechette.** 2009. “Collusion as Public Monitoring Becomes Noisy: Experimental Evidence.” *Journal of Economic Theory* 144(3): 1135–65.
- Bereby-Meyer, Yoella, and Alvin E. Roth.** 2006. “The Speed of Learning in Noisy Games: Partial Reinforcement and the Sustainability of Cooperation.” *American Economic Review* 96(4): 1029–42.
- Bergin, James, and W. Bentley MacLeod.** 1993. “Continuous Time Repeated Games.” *International Economic Review* 34(1): 21–37.
- Borenstein, Severin.** 2003. “Rapid Communication and Price Fixing: The Airline Tariff Publishing Company Case.” In *The Antitrust Revolution: The Role of Economics*, 4th ed., edited by John E. Kwoka and Lawrence J. White. New York: Oxford University Press.

- Cooper, Russell, Douglas V. DeJong, Robert Forsythe, and Thomas W. Ross.** 1996. "Cooperation without Reputation: Experimental Evidence from Prisoner's Dilemma Games." *Games and Economic Behavior* 12(2): 187–218.
- Dal Bo, Pedro.** 2005. "Cooperation under the Shadow of the Future: Experimental Evidence from Infinitely Repeated Games." *American Economic Review* 95(5): 1591–1604.
- Dal Bo, Pedro, and Guillaume R. Frechette.** 2011. "The Evolution of Cooperation in Infinitely Repeated Games: Experimental Evidence." *American Economic Review* 101(1): 411–29.
- Engle-Warnick, Jim, and Robert L. Slonim.** 2006. "Inferring Repeated-Game Strategies from Actions: Evidence from Trust Game Experiments." *Economic Theory* 28(3): 603–32.
- Flood, Merrill M.** 1952. "Some Experimental Games." RAND Research Memorandum RM-789-1.
- Fudenberg, Drew, and Eric Maskin.** 1986. "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information." *Econometrica* 54(3): 533–54.
- Hauk, Esther, and Rosemarie Nagel.** 2001. "Choice of Partners in Multiple Two-Person Prisoner's Dilemma Games: An Experimental Study." *Journal of Conflict Resolution* 45(6): 770–93.
- Hogg, Robert V. and Allen T. Craig.** 1970. *Introduction to Mathematical Statistics*, 3rd ed. New York: Macmillan.
- Huberman, Bernardo A. and Natalie S. Glance.** 1993. "Evolutionary Games and Computer Simulations." *Proceedings of the National Academy of Sciences* 90(16): 7716–18.
- Klein, Joel I.** 1998. "The State of Competition in the Airline Industry," Department of Justice Statement before the Committee on the Judiciary, United States House of Representatives.
- Kreps, David M., Paul Milgrom, John Roberts, and Robert Wilson.** 1982. "Rational Cooperation in the Finitely Repeated Prisoners' Dilemma." *Journal of Economic Theory* 27(2): 245–52.
- McKelvey, Richard D., and Thomas R. Palfrey.** 1995. "Quantal Response Equilibria for Normal Form Games." *Games and Economic Behavior* 10(1): 6–38.
- Radner, Roy.** 1986. "Can Bounded Rationality Resolve the Prisoner's Dilemma?" In *Contributions to Mathematical Economics*, edited by Andreu Mas-Colell and Werner Hildenbrand, 387–99. Amsterdam: North-Holland.
- Rapoport, Anatol, and Albert M. Chammah.** 1965. *Prisoner's Dilemma: A Study in Conflict and Cooperation*. Ann Arbor: University of Michigan Press.
- Rapoport, Anatol, and Melvin J. Guyer, and David G. Gordon.** 1976. *The 2x2 Game*. Ann Arbor: University of Michigan Press.
- Roth, Alvin E. and J. Keith Murnighan.** 1978. "Equilibrium Behavior and Repeated Play of the Prisoner's Dilemma." *Journal Mathematical Psychology* 17(2): 189–98.
- Selten, Reinhard, and Rolf Stoecker.** 1986. "End Behavior in Sequences of Finite Prisoner's Dilemma Supergames: A Learning Theory Approach." *Journal of Economic Behavior and Organization* 7(1): 47–70.
- Simon, Leo K., and Maxwell B. Stinchcombe.** 1989. "Extensive Form Games in Continuous Time: Pure Strategies." *Econometrica* 57(5): 1171–1214.