

Steady-State Equilibria in Anonymous Repeated Games, II: Coordination-Proof Strategies in the Prisoner's Dilemma*

Daniel Clark, Drew Fudenberg, and Alexander Wolitzky

October 25, 2019

Abstract

We study when and how cooperation can be supported in the repeated prisoner's dilemma in a large population with random matching and overlapping generations, when players have only *first-order information* about their current partners: a player's record tracks information about their past actions only, and not their partners' past actions (or their partners' partners' actions). We restrict attention to strict equilibria that are *coordination-proof*, meaning that two matched players never play a Pareto-dominated Nash equilibrium in the stage game induced by their records and expected continuation payoffs. We find that simple strategies can support limit efficiency if the stage game is either "mild" or "strongly supermodular," and that no cooperation can occur in equilibrium for a near-complementary parameter set. The presence of "supercooperator" records, where a player cooperates against any opponent, is crucial for supporting maximal cooperation when the stage game is "severe."

*We thank Nageeb Ali, V Bhaskar, Glenn Ellison, Sander Heinsalu, Yuval Heller, Takuo Sugaya, Satoru Takahashi, and Caroline Thomas for helpful comments and conversations, and NSF grants SES 1643517 and 1555071 and Sloan Foundation grant 2017-9633 for financial support.

1 Introduction

People often cooperate with sequences of different partners with little information about their partners' past behavior and almost no information about the past behavior of their partners' partners.¹ To study such settings, this paper analyzes the enforcement of cooperation when a large population of agents is randomly and anonymously matched to play the prisoner's dilemma with fairly minimal information. Our goals are to characterize how payoff parameters determine the possibility and maximum extent of "robust" cooperation, as well as what types of strategies are likely to be effective.

We use the state-state model of Clark, Fudenberg, and Wolitzky (2019b) ("CFW"), where there is a unit mass of agents with geometrically distributed lifetimes, pairs of agents are anonymously matched to play a stage game, and there is neither a known start date nor a sense of common calendar. Relative to CFW, here we focus on the prisoner's dilemma (rather than considering more stage games) and consider only *first-order* records, meaning that a player's information about their partner depends only on the actions the partner has taken, and not on the actions or histories of the partner's past partners. Within this more restricted environment, we provide necessary and sufficient conditions for cooperation under a much broader class of strategies than the trigger strategies studied in CFW.

We impose only two restrictions. First, we consider only strict equilibria; this captures a simple form of robustness and, in particular, rules out "belief-free" equilibria and related constructions. Second, we say an equilibrium is *coordination-proof* if whenever the "augmented" game faced by a pair of matched players (which depends on their current records and their expected continuation payoffs to each action) is a coordination game, they play the Pareto-dominant equilibrium. Restricting attention to coordination-proof strategies rules out equilibria built on within-match miscoordi-

¹Seminal studies of such "community enforcement" include Ellickson (1991), Friedman and Resnick (2001b), Greif (1993), Klein and Leffler (1981), Milgrom, North, and Weingast (1990), and Ostrom (1990). Clark, Fudenberg, and Wolitzky (2019b) contains additional references.

nation.

We characterize payoff parameters for which there can never be strict and coordination-proof equilibria with any cooperation. We then show that for a complementary set of parameters, equilibria with arbitrarily high levels of cooperation are possible when there is little noise and the players' expected lifetimes are long.

Recall the standard normalization of the prisoner's dilemma payoff matrix, where $g, l > 0$ and $g < l + 1$, so (C, C) maximizes the sum of payoffs.

	C	D
C	$1, 1$	$-l, 1 + g$
D	$1 + g, -l$	$0, 0$

Figure 1: The Prisoner's Dilemma

Here g measures the gain to defection (that is, playing D) when one's opponent cooperates, for example the cost savings from providing a low quality product or service, or the profit gained by cheating a business partner. Because l measures the gain from playing D against D , the comparison of g and l reflects the degree of complementarity, which we will see is an important factor in determining the possibility and maximal extent of equilibrium cooperation. Intuitively, this is because first-order information is not sufficient to distinguish between opportunistic deviations to D and equilibrium plays of D that punish opponents with bad records, so players must sometimes be willing to worsen their record by playing D against D when their continuation payoff would be higher if they played C and incurred a short run loss.

To help organize our results, we say that the game is *submodular* when $g \geq l$, *strictly supermodular* when $g < l$, and *strongly supermodular* when $g + g^2 < l$. We also say it is *mild* when $g < 1$ and *severe* when $g \geq 1$.

In our model, every strict equilibrium is symmetric, and the steady state where everyone always plays D regardless of the records is always a strict equilibrium. As in the related random matching models of Takahashi (2010) and Heller and Mohlin (2018), we find that when the prisoner's dilemma stage game is submodular, the only strict equilibrium is *Always Defect*, regardless of the community's (first-order) record

system. We thus focus on the supermodular case.

As a preliminary result, we show that in any strict equilibrium, each record is either a *defector* that defects regardless of the partner’s record, a *supercooperator* that cooperates regardless of the partner’s record, or a *preciprocator* that cooperates iff the equilibrium strategy says that the partner will cooperate with them. The presence of preciprocator records is what can provide incentives to cooperate. The presence of supercooperator records can make it easier to satisfy the constraint that preciprocators play D against defectors and can also increase steady-state cooperation by reducing the rate at which new plays of D enter the system.

Note that while supercooperator records cooperate with all partners, and preciprocators cooperate with all supercooperators, not all preciprocators records need cooperate with each other. When they do not, the induced one-shot game between two preciprocators is a coordination game with two strict equilibria. Coordination-proofness thus requires that two matched preciprocators always play the Pareto-dominant (C, C) equilibrium. This simplifies the analysis, as it implies that a strategy profile is completely characterized by a description of which records are preciprocators, which are supercooperators, and which are defectors.

Our main result, Theorem 3, provides necessary and sufficient conditions for cooperation in strict, coordination-proof equilibria. The necessary conditions (Theorem 3(a)) establish two facts. First, when $g \geq 1$, equilibrium cooperation requires the presence of supercooperator records. This is because without supercooperators only (C, C) and (D, D) are played on-path, so incentives require that each play of D leads to g “switches” from (C, C) to (D, D) . When $g \geq 1$, these switches “snowball,” precluding positive steady-state cooperation. Second, supercooperator records can exist only when $l > g + g^2$. Intuitively, higher values of g make it harder to support cooperation because deterring defection requires harsher punishments when g is larger, and these punishments occur with positive frequency on path. In contrast, higher values of l help support cooperation by preventing “undesired” cooperation against defectors.² The

²This force is similar to that in Kandori (1992), but Kandori’s construction requires deterring all

threshold $l = g + g^2$ for the emergence of supercooperation comes from combining such incentive effects with the equations for steady-state population shares.

The near-converse, Theorem 3(b), shows that full efficiency can be attained in the iterated limit where first lifespans go to ∞ and then noise goes to 0 when the game is supermodular and either mild or satisfies $\max\{g + g^2, f(g)\} < l$, where the function f is a nuisance term discussed below. When the game is mild, our proof uses *cyclic strategies* of the form $(P_K D_M)^\infty$, meaning that a record k is a reciprocator if $k \bmod (K+M) < K$ and a defector if $k \bmod (K+M) \geq K$. For example, $(P_1 D_1)^\infty$, or “even-odd,” is the cyclic strategy where even-numbered records reciprocate and odd-numbered records defect. With these strategies the incentive constraints of the defectors are satisfied for any parameters, because playing D simultaneously maximizes their current payoff and their continuation value. When $g < 1 < l$, noise is low, and expected lifetimes are long, even-odd satisfies the incentive constraints of the reciprocators and attains limit efficiency. When $g < l < 1$, even-odd is not an equilibrium in the limit, because the punishment for playing D (a loss of 1 next period) is so harsh that reciprocators would prefer to cooperate against defectors. In such cases we use other values of K and M to ensure that reciprocators are willing to play C against each other while playing D against defectors.

In severe prisoner’s dilemmas, strategies of the form $(P_K D_M)^\infty$ cannot support cooperation, because they do not have any supercooperators. Limit efficiency can be attained with cyclic strategies that also have supercooperators, but our proof instead uses *generalized trigger strategies* of the form $P_K S_L D_\infty$: here records $0 \leq k < K$ are reciprocators, $K \leq k < K + L$ are supercooperators, and $k \geq K + L$ are defectors. We use generalized trigger strategies because they are somewhat simpler, and also because as we show in Section 7 they are the only strategies that satisfy the additional robustness criterion of “forgery-proofness.”

The special case of generalized trigger strategies where $K = 1$ and $L = 0$ is the usual *Grim* strategy; the case with $K > 1$ and $L = 0$ corresponds to what Fudenberg,

supercooperation, while our construction requires a positive level of supercooperation on path.

Rand, and Dreber (2012) and Clark, Fudenberg, and Wolitzky (2019a) call *GrimK*.³ Trigger strategies with $L = 0$ cannot support cooperation when $g > 1$ but generalized trigger strategies with supercooperators can. However, as in noisy repeated games with fixed partners, any fixed generalized trigger strategy yields very low payoffs in the limit as $\gamma \rightarrow 1$, because in the resulting steady state most players will be defectors. Instead, we construct sequences of equilibria where the expected fraction of a player’s lifetime spent as a reciprocator is roughly constant as $\gamma \rightarrow 1$.

When $g \geq 1$ and $g + g^2 < l$, our construction requires the additional condition $f(g) < l$, so it is not quite a converse to the necessary conditions. The function f is a nuisance term that comes from the requirement that the records be integers, and as Theorem 4 shows the constraint $f(g) < l$ is not necessary with a richer record structure that circumvents this integer problem. Here we use “personal public randomizations,” which are a “decentralized” form of public randomization. This not only lets us dispense with the function f , but also allows the construction of equilibria with efficient payoffs for the general $(\gamma, \varepsilon) \rightarrow (1, 0)$ limit (rather than only a specific iterated limit).

Finally, we show that when the record system takes the form of a noisy count of how many times the player intended to play D , generalized trigger strategies are the only strategies that are both coordination-proof and satisfy the property that a player never benefits from “forging” records of additional interactions (somewhat like posting fake reviews on an online rating platform). For example, cyclic strategies are not forgery-proof, because any player with a defector record would forge enough additional D ’s to reach the next reciprocator record.

³For $K > 1$, *GrimK* is never an equilibrium in a two player game with perfect monitoring. Fudenberg, Rand, and Dreber (2012) note that it can be an equilibrium when actions are observed with noise, and that some experimental subjects seem to use such strategies. In our framework, the *GrimK* strategy class can never achieve limit efficiency. Clark, Fudenberg, and Wolitzky (2019a) analyze its limit performance and discuss its implications for indirect reciprocity.

1.1 Related Work

Rosenthal (1979) and Rosenthal and Landau (1979) introduced the study of repeated games with random matching. Rosenthal (1979) supposed that players know only the action that their current opponent played in the previous period, and showed that in the prisoner’s dilemma cooperation can be supported by pure strategy equilibria only for a particular knife-edge value of the discount factor.

Kandori (1992) and Ellison (1994) showed that cooperation in the prisoner’s dilemma can be enforced by “contagion equilibria” in finite populations even when the players have “zero-order” information—that is, no information beyond the play in their own matches—but the required discount factor converges to 1 as the population becomes infinitely large. Kandori constructed simple contagion equilibria that exist only under a fairly strong restriction on the payoff functions: the loss parameter l needs to be sufficiently large, and in particular must diverge to ∞ as $\delta \rightarrow 1$. Ellison extended Kandori’s results both to arbitrary payoff parameters and to approximately efficient equilibria in settings with a small amount of noise by using either public randomizing devices or “threading,” which both serve to lower the players’ effective discount factor, and so make the prescribed equilibrium punishments incentive-compatible.⁴ In noisy environments, Ellison’s “threads” construction yields equilibria with high initial expected payoffs even though play in each thread eventually converges to everyone playing D . In contrast, efficiency in our steady-state model requires that the overall cooperation rate never falls; in this sense our equilibria are what Kandori (1992) called “globally stable.”⁵

Three previous papers have studied continuum-player repeated games with anonymous random matching and first-order information: Takahashi (2010), Heller and

⁴Deb, Sugaya, and Wolitzky (2018) generalize these results to prove the folk theorem for finite-population repeated games with anonymous random matching. Their construction relies heavily on a finite population, common calendar time, and non-strict incentives.

⁵Moreover, in our continuum population the aggregate effect of noise is deterministic; this prevents Ellison’s argument that when the noise level is small, players who have not seen a play of D think that it is unlikely that any D ’s have been played.

Mohlin (2018), and Bhaskar and Thomas (2018).

Takahashi (2010) assumes players know the entire record of each partner’s past play—all “first-order” information. He shows that cooperation can be supported using somewhat complex and unintuitive “belief-free” mixed strategies, which are ruled out by our restriction to strict equilibria. He also shows that cooperation can be supported in strict equilibria when the players are patient and the game is strictly supermodular, provided there is no noise in the implementation or recording of actions. To do this, he follows Ellison and uses threads to calibrate the effective discount factor to be within the interval where players want to cooperate against a partner who cooperates and defect against a partner who defects. Because our game does not have a commonly known start date, such threads cannot be constructed here. Conversely, Takahashi shows that no strategies support cooperation as a strict equilibrium when payoffs are strictly submodular. The intuition is simple: with only first-order information, a player’s continuation payoff depends on their record and action today but not on their current partner’s action. Thus, in order for a player to strictly prefer to cooperate with a cooperating partner while defecting against a defecting partner, payoffs must be strictly supermodular, and moreover the effective discount factor must be low enough that the difference in payoffs caused by the partner’s current action can offset the difference in future continuation payoffs.

Heller and Mohlin (2018)’s study of the prisoner’s dilemma with anonymous random matching assumes that a small fraction of players are commitment types (some of whom are committed to strictly mixed actions), and that players live forever, are infinitely patient, and see only a finite sample of their partners’ past actions. Players are restricted to use stationary strategies that condition only on the sampled actions of their partners and not on their own histories. Here, as in Takahashi, there are no cooperative equilibria when payoffs are submodular, and for much the same reason: Players will play C or D depending on how likely it is their opponent plays C, so when their observation of their partner’s play consists only of D’s they are most likely to play C, but then there is no incentive to play C, and cooperation breaks down.

Conversely, Heller and Mohlin show that when payoffs are supermodular, the presence of commitment types allows the construction of an efficient and relatively simple mixed-strategy equilibrium without threads or public randomization, because a partner’s past actions are a noisy signal of their type and thus of their likely current-period action.⁶

Bhaskar and Thomas (2018) study a sequential-move “lending game” with one-sided moral hazard, where borrowers are constrained to default with a fixed i.i.d. probability. They show that cooperation can be supported if lenders are told only whether or not a borrower has defaulted in any of the last M periods for some sufficiently large M . The distinction between submodular and supermodular games does not arise here due to the sequential nature of the game.

Nowak and Sigmund (1998) and many subsequent papers study the enforcement of cooperation using “image scoring,” which means that each player has first-order information about their partner, but conditions their action only on their partner’s record and not their own record. These strategies are never a strict equilibrium for any stage game, and are typically unstable in environments with noise (Panchanathan and Boyd, 2003). One interpretation of our model is that it shows that image scoring-type strategies can be strict equilibria, provided the game is supermodular and players condition on their own record as well as their partner’s.

Finally, our coordination-proofness assumption is somewhat reminiscent of renegotiation-proofness in fixed-partner repeated games as studied by Farrell and Maskin (1989) and others, but it is simpler since each pair of partners plays only a one-shot game.⁷

⁶Heller and Mohlin also consider alternative information structures where players observe, for example, a finite sample of their partners’ past action *profiles*. Dilmé (2016) also assumes commitment types, and uses them to construct a belief-free cooperative equilibrium for the case where $g = l$.

⁷Our interest in “simple” strategies is also motivated by laboratory studies of repeated games suggesting that many subjects use fairly simple strategies, e.g. Fudenberg, Rand, and Dreber (2012), Dal Bó and Fréchette (2018).

2 Records, Steady States, and Equilibria

This section introduces the class of record system we will consider along with our matching model, and briefly reviews some relevant definitions and results from CFW.

Throughout, we consider a discrete time model with a constant unit mass of players, each of whom has a geometrically-distributed lifespan with continuation probability $\gamma \in (0, 1)$, with exits balanced by a steady flow of new entrants. The time horizon is doubly infinite. When a pair of players match, they play the prisoner’s dilemma with action set $A = \{C, D\}$ and payoffs given by Figure 1.

2.1 First-Order Records

We suppose that when players match they observe each other’s record, which is an element r of a countable set R , but no other information. New players all enter with the same initial record, which we denote by 0.⁸

In this paper we consider only first-order records, meaning that each player’s record is updated (possibly stochastically) based only on their own action.⁹

Definition 1. *A **first-order record system** is a function $\rho : R \times A \rightarrow \Delta(R)$ that specifies a probability distribution over a player’s next-period record given the player’s current-period record r and current-period action a .*

There are at least three reasons the record system might be stochastic: First, randomization may arise by design, as it makes it easier to satisfy the various incentive constraints required to enforce equilibrium cooperation. Second, randomization may be due to some unavoidable probability of a “recording error.” And third, there may be “implementation errors,” so that a player who intended to play C plays D instead.¹⁰

⁸Our main results, Theorems 3 and 4, extend to the case where initial records have finite support.

⁹See CFW for an analysis of second-order record systems, where a player’s updated record can depend on both players’ actions.

¹⁰Under the implementation errors interpretation, the stage-game payoff matrix represents the expected values of the underlying payoffs, and we require that the outcome of the game has a product structure in the sense of Fudenberg, Levine, and Maskin (1994), so that the updating of a player’s record does not depend on their opponent’s action.

The following definition applies under the recording error or implementation error interpretation.

Definition 2. A first-order record system is **noisy** if for each record r there exist $q_C(r), q_D(r) \in \Delta(R)$ and $\nu_C(r) \in (0, 1/2]$, $\nu_D(r) \in [0, 1/2]$ such that

- $\rho(r, C) = (1 - \nu_C(r))q_C(r) + \nu_C(r)q_D(r)$,
- $\rho(r, D) = \nu_D(r)q_C(r) + (1 - \nu_D(r))q_D(r)$.

Here $q_C(r)$ represents the distribution over records after “a recording of C is fed into the record system,” $q_D(r)$ represents the distribution over records after “a recording of D is fed into the record system,” and the ν ’s represent the noise, which can either be in the recording system itself (in the “recording errors” interpretation) or in the map from intended to realized actions (the “implementation errors” interpretation). Our conditions for limit efficiency apply when there is some chance of recording or implementation errors, as reflected in the assumption that $\nu_C(r) > 0$.

The following simple noisy record system will play an important role in our analysis:

Definition 3. Noisy Counting D ’s is a noisy first-order record system where $R = \mathbb{N}$, new players have record 0, $\rho(k, C) = (1 - \varepsilon)\delta_k + \varepsilon\delta_{k+1}$ for some $\varepsilon > 0$, and $\rho(k, D) = \delta_{k+1}$, where δ_k denotes a degenerate distribution on record k .¹¹

We will see that, when noise is small and lifespans are long, Noisy Counting D ’s supports full efficiency for any parameters under which any cooperation is possible for any noisy first-order record system.¹² Since partners can always agree not to condition their behavior on jointly observed random variables, the same result holds a fortiori for the following richer class of record systems:

¹¹We will typically use k to denote a player’s record in the Noisy Counting D ’s record system, whereas we typically use r in general record systems. We do this in part to make it clear when we are analyzing the Noisy Counting D ’s record system in particular, but also because k will sometimes reflect a “score” that is derived from the player’s record, rather than the record itself, when we consider record systems that are “richer than Noisy Counting D ’s” in the sense we define shortly.

¹²This result extends to the case where plays of D can also be mis-recorded as C ; we rule this out to simplify some formulas.

Definition 4. A first-order record system is *richer than Noisy Counting D 's* if $R = \mathbb{N} \times R'$ for some countable set R' and the projection of ρ on its first component coincides with Noisy Counting D 's for some $\varepsilon > 0$.

For example, the record system that keeps separate noisy counts of the number of times a player has played D and C is richer than Noisy Counting D 's, and so is the record system that noisily counts D 's while also recording the player's age at each period where they played D . Note that the Noisy Counting D 's record system is indexed by a single parameter ε , and any record system that is richer than Noisy Counting D 's can be partially indexed by the corresponding ε .

2.2 Strategies and Steady States

Since there is a continuum of players, only a player's current record and that of their current partner matter for the player's current payoff, and only the player's own record will matter in the future. For this reason, all strict equilibria are *pairwise-public*, meaning that they condition only on information that is public knowledge between the two partners, namely their records. We write a pairwise-public pure strategy as a function $\mathbf{s} : R \times R \rightarrow A$, with the convention that the first coordinate is the player's own record and the second coordinate is that of the partner. Since we will restrict attention to strict equilibria, we consider only pairwise-public pure strategies. Moreover, since every strict equilibrium in a symmetric, continuum-population model is symmetric, we also assume all players use the same strategy.

The *state* of the system is the share of players with each possible record; we denote this by $\mu \in \Delta(R)$. We model random matching as in CFW: when the current state is μ , the distribution of matches is given by $\mu \times \mu$, so that, for each $(r, r') \in R^2$ with $r \neq r'$, the fraction of matches between players with record r and r' is $2\mu_r\mu_{r'}$, while the fraction of matches between two players with record r is μ_r^2 .

Given a record system and a strategy \mathbf{s} , we can define the *update map* from current states to next-period states, $f_{\mathbf{s}} : \Delta(\mathcal{R}) \rightarrow \Delta(\mathcal{R})$, as in CFW. A *steady state* under \mathbf{s} is

a state μ such that $f_{\mathbf{s}}(\mu) = \mu$.

Theorem 1. (CFW) *Under any first-order record system and any pairwise-public strategy, a steady state exists.*

Note that Theorem 1 does not assert that the steady state for a given strategy is unique, and indeed CFW give an example where it is not. Intuitively, this multiplicity corresponds to different initial conditions at time $t = -\infty$.

It remains to define equilibrium. Given a strategy \mathbf{s} and state μ , define the flow payoff of a player with record r as

$$\pi_r(\mathbf{s}, \mu) = \sum_{r'} \mu_{r'} u(\mathbf{s}(r, r'), \mathbf{s}(r', r)).$$

Next, denote the probability that a player with record r today has record r' t periods from now by $\phi_r(\mathbf{s}, \mu)^t(r')$. The continuation value of a player with record r is then given by

$$V_r(\mathbf{s}, \mu) = (1 - \gamma) \sum_{t=0}^{\infty} \gamma^t \sum_{r'} \phi_r(\mathbf{s}, \mu)^t(r') \pi_{r'}(\mathbf{s}, \mu).$$

Note that we have normalized continuation payoffs by $(1 - \gamma)$ to express them in per-period terms.

A pair (\mathbf{s}, μ) is a *strict equilibrium* if μ is a steady state under \mathbf{s} and, for each own record r and opponent's record r' , we have

$$\mathbf{s}(r, r') = \arg \max_{a \in A} \left[(1 - \gamma) u(a, \mathbf{s}(r', r)) + \gamma \sum_{r''} (\rho(r, a)[r'']) V_{r''}(\mathbf{s}, \mu) \right].$$

Thus, a player's objective is to maximize their expected undiscounted lifetime payoff. The existence of strict equilibria follows immediately from Theorem 1.

Corollary 1. *Under any first-order record system, a strict equilibrium exists. In particular, Always Defect ($\mathbf{s}(r, r') = D$ for all (r, r')) is always a strict equilibrium strategy.*

Proof. If $\mathbf{s}(r, r') = D$ for all (r, r') then (\mathbf{s}, μ) is a strict equilibrium for any steady state μ . ■

3 Preliminary Results

We introduce some basic concepts that are used repeatedly in our analysis.

3.1 Defectors, Supercooperators, Preciprocators

Definition 5. *Given an equilibrium (\mathbf{s}, μ) , record r is a*

- **defector** if $\mathbf{s}(r, r') = D$ for all r' .
- **supercooperator** if $\mathbf{s}(r, r') = C$ for all r' .
- **preciprocator** if $\mathbf{s}(r, r') = \mathbf{s}(r', r)$ for all r' , and moreover there exist r', r'' such that $\mathbf{s}(r, r') = C$ and $\mathbf{s}(r, r'') = D$.

Defectors play D against all partners, while supercooperators play C against all partners, even those who will play D against them. In contrast, preciprocator records exhibit a form of anticipatory reciprocation: they play C with partners they expect to play C , but play D with partners they expect to play D . We show in Section 5.2 that cooperation in coordination-proof equilibria requires some supercooperator records whenever $g \geq 1$.

Recall that the prisoner's dilemma is strictly supermodular if $g < l$, so the benefit of defecting is greater when the opponent defects. Conversely, the stage game is strictly submodular when $g > l$. A leading example of the prisoner's dilemma is reciprocal gift-giving, where each player can pay a cost $c > 0$ to give their partner a benefit $b > c$. In this case, a player receives the same static gain from playing D instead of C regardless of the play of their opponent, so $g = l$, and the game is neither strictly supermodular nor strictly submodular. Bertrand competition (with two price levels $H > L$) is supermodular whenever $L > H/2$ (the condition for the game to be a prisoner's dilemma), and Cournot competition (with two quantity levels) is submodular whenever marginal revenue is decreasing in the opponent's quantity.

Lemma 1. *In any strict equilibrium:*

1. If $g \geq l$ then every record is a defector or a supercooperator.
2. If $g < l$ then every record is a defector, supercooperator, or reciprocator.

Proof. Fix a strict equilibrium. Because the records use only first-order information, each player's continuation payoff depends only on their current record and action, so the optimal action in each period depends only on the player's record and the action prescribed by their opponent's record.

Suppose that $g \geq l$. Fix a record r , and suppose two players who both have record r meet. By symmetry, they play either (C, C) or (D, D) . In the former case, C is the strict best response to C . Since the current-period gain from playing D instead of C is weakly smaller when the opponent plays D , this means C is also the strict best response to D , so record r is a supercooperator. In the latter case, D is the strict best response to D , and therefore is also the strict best response to C , so record r is a defector.

If $g < l$ and D is strictly optimal against C , then D is also strictly optimal against D , so every record is either a defector, a supercooperator, or a reciprocator. ■

Theorem 2. *If $g \geq l$, the unique strict equilibrium is Always Defect.*¹³

Proof. By Lemma 1, if $g \geq l$ then the distribution of opposing actions faced by any player is independent of their record, so D is always optimal. ■

In what follows, we restrict to the strictly supermodular case where $g < l$. Here it is possible for some records to be reciprocators, which is what will allow equilibria that support some cooperation.

¹³Takahashi (2010) and Heller and Mohlin (2018) obtain the same conclusion in related models. This necessary condition applies for any first-order record system, but not for second-order records (CFW), or if the players have access to correlating devices, as then Lemma 1 fails because players with the same record could randomize between (C, D) and (D, C) . However, even with correlating devices the amount of cooperation is bounded away from 1 if $g \geq l$. The conclusion of Theorem 2 also extends to (possibly non-strict) pure-strategy equilibria whenever $g > l$.

3.2 Coordination-Proofness

Coordination-proofness is based on the idea that equilibria that rely on “miscoordination” within a match will fall apart if matched partners manage to coordinate successfully.

For a fixed steady-state equilibrium, denote the expected continuation payoff of a player with record r who plays action a by $V_r^a := \sum_{r'} \rho(r, a)[r']V_{r'}$. The *augmented payoffs* are then $\hat{u}_r(a, a') := (1 - \gamma)u(a, a') + \gamma V_r^a$. The *augmented game* is the static game with augmented payoffs, as shown in Figure 2.

	C	D
C	$(1 - \gamma)1 + \gamma V_r^C$	$-(1 - \gamma)l + \gamma V_r^C$
D	$(1 - \gamma)(1 + g) + \gamma V_r^D$	$0 + \gamma V_r^D$

Figure 2: The payoffs of a player with record r in the augmented stage game.

By definition, preciprocators play C against opponents who play C and play D against those who play D , so the augmented stage game between any two preciprocators is a coordination game where both (C, C) and (D, D) are stage-game Nash equilibria. Since playing D always gives a short-run gain, the fact that preciprocators play C against C implies that cooperating leads to higher continuation payoffs, so the (C, C) equilibrium yields both higher continuation payoffs and higher stage-game payoffs than the equilibrium where both play D . This observation motivates the following definition:

Definition 6. *An equilibrium is **coordination-proof** if whenever two preciprocators match, they play (C, C) .*¹⁴

Coordination-proofness implies that every preciprocator plays C when matched with another preciprocator or a supercooperator, and plays D when matched with a defector, so all preciprocators play C against the same set of opposing records. Thus a strategy profile is completely characterized by a description of which records

¹⁴Because two matched players will never face each other again, the coordination problem here is much simpler than that in dynamic games with a fixed set of partners, as in e.g. Bernheim and Ray (1989), Farrell and Maskin (1989), and Chassang and Takahashi (2011).

are reciprocators, which are supercooperators, and which are defectors. Denote the total population share in each class by μ^P , μ^S , and μ^D respectively. We will use the term *cooperator* for all players who are either reciprocators or supercooperators (i.e., anyone who is not a defector), and we denote the population share of cooperators by $\mu^C = \mu^P + \mu^S = 1 - \mu^D$.

4 Cooperation and Limit Efficiency

This section states our main result, which characterizes the parameters under which limit efficiency is attainable in coordination-proof equilibria with noisy first-order records. The following section sketches the proof.

Definition 7. *For any first-order record system that is richer than Noisy Counting D 's, **limit efficiency** is attainable if, for every $\eta > 0$, there exists $\bar{\varepsilon} < 1$ and a function $\bar{\gamma} : (0, 1) \rightarrow (0, 1)$ such that, whenever $\varepsilon < \bar{\varepsilon}$ and $\gamma > \bar{\gamma}(\varepsilon)$, there exists a coordination-proof strict equilibrium with $V_0 > 1 - \eta$.*

This iterated limit matches the order of limits in related papers on almost-perfect private monitoring (e.g., Ellison (1994) and Hörner and Olszewski (2006)). Note that V_0 , the per-period expected payoff of a newborn agent, is also equal to the average payoff in the population in every period. This follows because the expected fraction of a player's lifetime spent at record r is equal to the fraction of the population with record r (and there is no discounting, so both V_0 and the population-average payoff are given by $\sum_r \mu_r \pi_r$). Thus, when limit efficiency is attainable, the total population payoff approaches its maximum possible value in the iterated limit where first $\gamma \rightarrow 1$ and then $\varepsilon \rightarrow 0$.

Recall that the prisoner's dilemma is *mild* if $g < 1$ and *severe* otherwise, and that the game is *strongly supermodular* if $l > g + g^2$. Define the function $f : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ by $f(g) = \min \left\{ \frac{1+g}{|\ln(1+g)-1|}, 21.9223 - 3.57143g \right\}$.¹⁵ The following is our main result.

¹⁵21.9223 - 3.57143g is the approximate form of $(3e^\phi - 2 - 2g)/(\phi - 1)$ when $\phi = 1.56$.

Theorem 3.

- (a) *For any noisy record system, if $g \geq 1$ and $l \leq g + g^2$ (i.e., the prisoner's dilemma is severe and not strongly supermodular), the only coordination-proof equilibrium is Always Defect.*
- (b) *For any noisy record system that is richer than Noisy Counting D's, if either $g < 1$ or $l > \max\{g + g^2, f(g)\}$, limit efficiency is attainable.*

The proof uses the function f to rule out parameters for which we do not know how to construct sequences of limit efficient strategies given our restriction to finite-support updates. Note however that $f(g) < g + g^2$ whenever $g > 2.9$. Section 6 proves the following result, which shows that if we introduce real-valued *personal public randomizations* the function f is not needed. Moreover, personal public randomizations allow the share of cooperators to converge to 1 in the general $(\gamma, \varepsilon) \rightarrow (1, 0)$ limit.

Theorem 4. *When personal public randomizations are available,*

- (a) *For any noisy record system, if $g \geq 1$ and $l \leq g + g^2$, the only coordination-proof equilibrium is Always Defect.*
- (b) *For any noisy record system that is richer than Noisy Counting D's, if either $g < 1$ or $l > g + g^2$ (i.e., the prisoner's dilemma is either mild or strongly supermodular), the maximal coordination-proof strict equilibrium value of V_0 converges to 1 as $(\gamma, \varepsilon) \rightarrow (1, 0)$.*

Figure 3 displays the conclusions of Theorems 3 and 4. Note that as g increases from just below 1 to just above 1, the critical value of l above which cooperation is possible jumps from 1 to at least 2.

Section 1 discussed the intuition for why small values of g and large values of l make supporting cooperation easier. The specific necessary condition $g < 1$ or $l > g + g^2$ comes from combining two inequalities: $\mu^S < 1/(1 + g)$ and $\mu^P + \mu^S(l - g) > g$. The inequality $\mu^S < 1/(1 + g)$ guarantees that the flow payoff to a defector of $\mu^S(1 + g)$ is

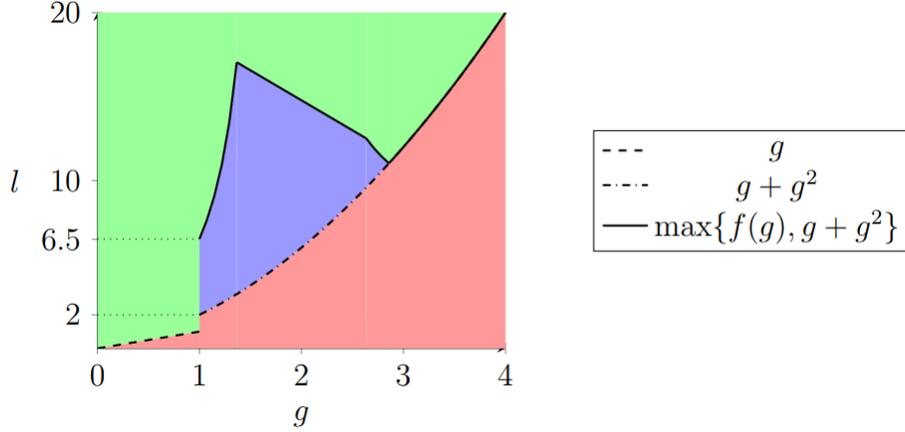


Figure 3: Limit efficiency obtains in the green region. In the blue region, limit efficiency obtains with PPR, but our results do not establish whether it obtains without PPR. In the red region, there are no cooperative equilibria.

less than 1, which is a necessary condition for an equilibrium with cooperation, since otherwise it would be optimal for newborn players to play D for their entire lives rather than following the equilibrium strategy. The inequality $\mu^P + \mu^S(l - g) > g$ is a necessary condition for cooperators to play C against C , as shown by Lemma 6 below. The next lemma shows that it is impossible to satisfy these two inequalities simultaneously when $g \geq 1$ and $l \leq g + g^2$.

Lemma 2. *It is not possible to satisfy both $\mu^S < 1/(1 + g)$ and $\mu^P + \mu^S(l - g) > g$ when $g \geq 1$ and $l \leq g + g^2$.*

Proof. Note the highest value of $\mu^P + \mu^S(l - g)$ is bounded above by either 1, which corresponds to $\mu^P = 1$ and $\mu^S = 0$, or $l/(1 + g)$, which corresponds to $\mu^P = g/(1 + g)$ and $\mu^S = 1/(1 + g)$. When $g \geq 1$, it must therefore be that $l/(1 + g) > g$, which requires $l > g + g^2$. ■

To see how we construct strategies that achieve limit efficiency when $g \geq 1$ and $l > g + g^2$, consider a strategy profile in which players begin with preciprocator records, and then over time transition to being supercooperators, and eventually to being defectors. Further assume that all players within a given class face exactly the same incentive

constraints, and that we can choose any steady-state shares of preciprocators and supercooperators, μ^P and μ^S , provided they sum to no more than 1. Both of these things will be true in the equilibria with personal public randomizations that we construct. When $l > g + g^2$, we can satisfy both the $\mu^S < 1/(1 + g)$ and $\mu^P + \mu^S(l - g) > g$ constraints. When ε is sufficiently small, the latter constraint implies that preciprocators prefer to play C against C . Additionally, when $l > g \geq 1$ preciprocators' incentives to play D against D are satisfied whenever supercooperators are willing to play C against D . We can also choose steady-state shares that satisfy $\mu^C - \mu^D l > \mu^S(1 + g)$, which guarantees that the flow payoff to a supercooperator is higher than that to a defector and ensures that supercooperators' incentives are satisfied when ε is sufficiently small. Thus, all of the incentive constraints can be satisfied for some μ^P and μ^S for small ε , and these shares can be chosen so that $\mu^C = \mu^P + \mu^S$ converges to 1 as ε goes to 0.

When $g < 1$, we can achieve limit efficiency in equilibria without supercooperators by targeting a sequence of preciprocator shares μ^P converging to 1 for which the incentives of preciprocators are always satisfied. Formally, this is because preciprocators' incentives are satisfied iff

$$g < \frac{(1 - \varepsilon)(1 - \mu^P)}{1 - (1 - \varepsilon)\mu^P} \mu^P < l,$$

and, for sufficiently small ε , these inequalities can be satisfied with μ^P converging to 1 as $\varepsilon \rightarrow 0$.

5 Proof of Theorem 3

5.1 Incentive Constraints with Noisy Records

We first derive the players' incentive constraints for an arbitrary noisy record system. Throughout, $(C|C)_r$ denotes the condition that C is the best response to C for a player with record r , $(C|D)_r$ denotes the condition that C is the best response to D , $(D|D)_r$

the condition that D is the best response to D , and $(C|D)_r$ the condition that C is the best response to D .

Let V_r^C denote the expected continuation payoff when a recording of C is fed into the record system for a record r player: that is, $V_r^C = E_{r' \sim q_C(r)}[V_{r'}]$, where $E_{r' \sim q_C(r)}$ indicates the expected value when r' is distributed according to $q_C(r)$. Similarly, let $V_r^D = E_{r' \sim q_D(r)}[V_{r'}]$ denote the expected continuation payoff when a recording of D is fed into the record system. Let π_r denote the expected flow payoff to a record r player under the equilibrium strategy, and let p_r^D denote the probability that a recording of D will be fed into the record system for a record k player. Note that $p_r^D > 0$ for all r since $\nu_C(r) > 0$ and $\nu_D(r) < 1$.

Given a noisy record system and an equilibrium, define the *normalized reward* for playing C rather than D for a record r player by

$$W_r := \frac{1 - \nu_C(r) - \nu_D(r)}{p_r^D} \left(\pi_r - V_r + \frac{\gamma}{1 - \gamma} (V_r^C - V_r) \right).$$

Lemma 3. *For any noisy record system,*

- *The $(C|C)_r$ constraint is $W_r > g$.*
- *The $(C|D)_r$ constraint is $W_r > l$.*
- *The $(D|D)_r$ constraint is $W_r < l$.*
- *The $(D|C)_r$ constraint is $W_r < g$.*

Proof. Consider a player with record r . We derive the $(C|C)_r$ constraint; the other constraints can be similarly derived. When a record r player plays C , their expected continuation payoff is $(1 - \nu_C(r))V_r^C + \nu_C(r)V_r^D$, since a recording of C is fed into the record system with probability $1 - \nu_C(r)$ and a recording of D is fed into the record system with probability $\nu_C(r)$. Similarly, when the player plays D , their expected continuation payoff is $\nu_D(r)V_r^C + (1 - \nu_D(r))V_r^D$. Thus, the $(C|C)_r$ constraint is $1 -$

$\gamma + \gamma(1 - \nu_C(r))V_r^C + \gamma\nu_C(r)V_r^D > (1 - \gamma)(1 + g) + \gamma\nu_D(r)V_r^C + (1 - \nu_D(r))V_r^D$, which is equivalent to

$$(1 - \nu_C(r) - \nu_D(r))\frac{\gamma}{1 - \gamma}(V_r^C - V_r^D) > g.$$

Note that $V_r = (1 - \gamma)\pi_r + \gamma(1 - p_r^D)V_r^C + \gamma p_r^D V_r^D$. Manipulating this gives $V_r^C - V_r^D = ((1 - \gamma)\pi_r - V_r + \gamma V_r^C)/(\gamma p_r^D)$. Substituting this into the above inequality gives the desired form of the $(C|C)_r$ constraint. ■

The incentive constraints take a simpler form in the special case of the Noisy Counting D 's record system. With Noisy Counting D 's, for all k we have $\nu_C(k) = \varepsilon$, $\nu_D(k) = 0$, $V_k^C = V_k$, and $V_k^D = V_{k+1}$. The normalized reward thus simplifies to

$$W_k = \frac{1 - \varepsilon}{p_k^D}(\pi_k - V_k).$$

Lemma 4. *For Noisy Counting D 's, Lemma 3 holds with $W_k = (1 - \varepsilon)(\pi_k - V_k)/p_k^D$.*

5.2 Necessary Conditions for Cooperation

Lemma 5. *In any equilibrium, $\mu^S < 1/(1 + g)$.*

Proof. Because new players with record 0 have the option of always playing D , in any equilibrium with $\mu^C > 0$, it must be that $\mu^S(1 + g) < V_0 \leq 1$, which gives $\mu^S < 1/(1 + g)$. ■

This lemma says that there cannot be too many supercooperators, as otherwise defectors would achieve the highest payoffs. Conversely, the next lemma implies that cooperation requires a positive share of supercooperators when $g \geq 1$, and moreover that the required share grows when g and l are increased by the same amount.

Lemma 6. *In any coordination-proof equilibrium with $\mu^C > 0$, $\mu^P + \mu^S(l - g) > g$.*

The proof of Lemma 6 is in A.1. It uses Lemma 12, also in A.1, which shows that there must exist a reciprocator or supercooperator record r with continuation value

sufficiently above the average payoff in the population, and then combines this with the incentive constraints of Lemma 3 to show that $\mu^P + \mu^S(l - g) > g$ must hold in any cooperative equilibrium.

Theorem 3(a) follows from Lemmas 5 and 6, since from Lemma 2 it is impossible to satisfy both $\mu^S < 1/(1 + g)$ and $\mu^P + \mu^S(l - g) > g$ when $g \geq 1$ and $l \leq g + g^2$.

5.3 Sufficient Conditions for Limit Efficiency

We now prove Theorem 3(b), which shows that limit efficient outcomes exist for “most” parameters where the necessary conditions of Theorem 3(a) are satisfied. We restrict attention here to Noisy Counting D 's; the same result holds a fortiori for any richer record system.

5.3.1 Cyclic Strategies Support Limit Efficiency when $g < 1$

We first outline how cyclic strategies can attain limit efficiency when the prisoner's dilemma is mild.

Fix any rational number ρ satisfying $g < \rho < \min\{l, 1\}$. Let K and M be integers such that $K \geq M > 0$ and $M/K = \rho$. We will show that $(P_K D_M)^\infty$ achieves limit efficiency.

We first establish that under this strategy, for fixed ε , the share of cooperators μ^C converges to $\bar{\mu}^C(\varepsilon)$ as $\gamma \rightarrow 1$, where

$$\bar{\mu}^C(\varepsilon) = \frac{1 + \rho - \sqrt{(1 + \rho)^2 - 4(1 - \varepsilon)\rho}}{2(1 - \varepsilon)\rho}. \quad (1)$$

Lemma 7. *With cyclic strategies, let $\mu^C(\gamma, \varepsilon)$ denote the share of cooperators in some steady state for arbitrary $(\gamma, \varepsilon) \in (0, 1) \times (0, 1)$. Then $\lim_{\gamma \rightarrow 1} \mu^C(\gamma, \varepsilon) = \bar{\mu}^C(\varepsilon)$.*

We give a heuristic argument for this result, the proof of which is in A.2. As $\gamma \rightarrow 1$, the mass μ^C of preciprocators will be approximately equally distributed among the K preciprocator phases, the mass $1 - \mu^C$ of defectors will be approximately equally

distributed among the M defector phases, existing players almost never die, and almost no newborn players enter the system. Since the share of players in phase $K - 1$ is approximately μ^C/K , the flow from phase $K - 1$ to phase K (the “outflow from cooperation”) is approximately $(1 - (1 - \varepsilon)\mu^C)\mu^C/K$, while the flow from phase $K + M - 1$ to phase 0 (the “inflow into cooperation”) is approximately $(1 - \mu^C)/M$, since the share of players in phase $K + M - 1$ is approximately $(1 - \mu^C)/M$. Setting these flows to be equal and solving for μ^C gives $\mu^C = \bar{\mu}^C(\varepsilon)$. Moreover, $\lim_{\varepsilon \rightarrow 0} \bar{\mu}^C(\varepsilon) = 1$, so $\lim_{\varepsilon \rightarrow 0} \lim_{\gamma \rightarrow 1} \mu^C = 1$.

The next step is to show that these strategies are strict equilibria in the iterated limit. The strategy even-odd, which is $(P_1 D_1)^\infty$, is a strict equilibrium in the iterated limit for $g < 1 < l$. To see this, note that g , the one-shot gain to a preciprocator from defecting when their opponent is a preciprocator, is less than the expected next-period loss of approximately 1, while l , the one-shot gain to a preciprocator from defecting when their opponent is a defector, is greater than the expected next-period loss. Since a preciprocator’s record in two periods will be even again with high probability regardless of their play today, preciprocation is indeed optimal. If $g < l < 1$, then even-odd is not an equilibrium, because the one-period punishment for playing D is too harsh to satisfy the $(D|D)$ constraint, but with $(P_K D_M)^\infty$ each play of D leads to only $\rho \leq 1$ periods of punishment. Formally, we prove the following lemma, which when combined with Lemma 4 and the fact that $p_k^D = 1 - (1 - \varepsilon)\mu^C$ for any preciprocator record, shows that the $(P_K D_M)^\infty$ strategy gives a strict equilibrium in the iterated limit.

Lemma 8. *With cyclic strategies, let $\mu^C(\gamma, \varepsilon)$ denote the share of cooperators and let $V_i(\gamma, \varepsilon)$ denote the value function for phase i in some steady state for arbitrary $(\gamma, \varepsilon) \in (0, 1) \times (0, 1)$. Then,*

$$\lim_{\varepsilon \rightarrow 0} \lim_{\gamma \rightarrow 1} \frac{1 - \varepsilon}{1 - (1 - \varepsilon)\mu^C(\gamma, \varepsilon)} (\mu^C(\gamma, \varepsilon) - V_i(\gamma, \varepsilon)) = \rho$$

for all $0 \leq i \leq K - 1$.

Combining these two lemmas yields

Lemma 9. *Limit efficiency is attainable in strict equilibrium using cyclic strategies whenever the prisoner’s dilemma is mild ($g < 1$).*

5.3.2 Generalized Trigger Strategies Support Limit Efficiency when $l > \max\{g + g^2, f(g)\}$

We now show that there are ways of choosing K and L as functions of γ and ε so that the class $P_K S_L D_\infty$ of generalized trigger strategies attains limit efficiency.

Lemma 10. *When $l > \max\{g + g^2, f(g)\}$, limit efficiency is attainable in strict equilibrium using generalized trigger strategies.*

Lemma 10 says that when the prisoner’s dilemma is sufficiently supermodular, generalized trigger strategies achieve limit efficiency. Note that when $g + g^2 > f(g)$ (i.e., $g > \approx 2.858$), limit efficiency is achieved if $l > g + g^2$. For smaller values of g , we use the stronger requirement that $l > f(g)$ to guarantee limit efficiency.

We do not know whether the condition $l > f(g)$ is necessary; it comes from the combination of the record $K - 1$ preciprocator and record K supercooperator facing identical incentives as $\gamma \rightarrow 1$ and the fact that K and L must be integers. In particular, when γ is near 1, the cost of increasing one’s record by 1 is very similar for a player with the last preciprocator record $K - 1$ and a player with the first supercooperator record K , so only a very narrow range of population shares is consistent with the incentive constraints. This narrow range is a problem because K and L must be integers; not all population shares are feasible in a steady state. However, when $l > f(g)$, we show that suitable population shares can be supported under which the incentives of both records $K - 1$ and K are satisfied. Section 6 shows that if there are stochastic elements to the record system that can be “tuned” as desired with the parameters, these integer problems can be ignored and $l > g + g^2$ suffices for limit efficiency.

Section A.3 gives a detailed outline of the proof of Lemma 10. The proof identifies a target level of cooperation for fixed ε , shows that there are feasible profiles satisfying

the incentive constraints where the level of cooperation actually attains this target as $\gamma \rightarrow 1$, and establishes that the level of cooperation can be sent to 1 as $\varepsilon \rightarrow 0$.

6 Personal Public Randomizations

Theorem 3 has two limitations: It falls short of completely characterizing when limit efficiency is attainable due to the presence of the “integer problem” term $f(g)$, and it only shows that efficiency is attainable for the iterated limit where first $\gamma \rightarrow 1$ and then $\varepsilon \rightarrow 0$, so that most players live long enough that their records will be perturbed by many occurrences of noise. This iterated limit seems most relevant to us, but it is also interesting to consider the implications of other ways of taking the double limit, including cases where ε is small compared to $1 - \gamma$, so very few players encounter any noise at all. In this case, if players start out as cooperators and only become defectors due to noise, there may be so few defectors that the $(D|D)$ constraint fails.

In this section, we show that enriching the information structure to allow a simple class of randomizing devices, namely *personal public randomizations (PPR)*, circumvents the “integer problem” and also allows the construction of asymptotically efficient equilibria for any form of the double limit $(\gamma, \varepsilon) \rightarrow (1, 0)$. We introduced the concept of PPR in CFW, where they allowed the construction of simpler equilibria and also let us close the gap between necessary and sufficient conditions for the play of certain actions to be fully supported in the double limit $(\gamma, \varepsilon) \rightarrow (1, 0)$.

With PPR, whenever a play of D is recorded, this play is associated in the player’s record with an indelible random variable ω that is uniformly distributed on $[0, 1]$. The record of these PPR can then be used by the player’s future partners to determine whether or not the recorded D should lead to punishment. PPR are thus a “decentralized” form of public randomization.¹⁶

Using PPR we can prove Theorem 4, which we restate for ease of reference:

¹⁶By covering the two extreme cases of no correlating devices whatsoever and PPR, we can give a good sense what may be achievable in intermediate cases where cases players have access to some form of imperfect correlating device whose outcome is captured by their records.

Theorem 4. *When personal public randomizations are available,*

- (a) *For any noisy record system, if $g \geq 1$ and $l \leq g + g^2$, the only coordination-proof equilibrium is Always Defect.*
- (b) *For any noisy record system that is richer than Noisy Counting D's, if either $g < 1$ or $l > g + g^2$, the maximal coordination-proof strict equilibrium value of V_0 converges to 1 as $(\gamma, \varepsilon) \rightarrow (1, 0)$.*

The proof of Theorem 4(a) is the same as that of Theorem 3(a), as Lemmas 5 and 6 still apply when PPR are available. The $l > g + g^2$ part of Theorem 4(b) extends the generalized trigger strategy construction of Theorem 3(b) by constructing strategies that use PPR to govern the way that records move through four “phases” labelled $D1$, P , S , and $D2$, where P is a preciprocator phase, S is a supercooperator phase, and $D1$ and $D2$ are defector phases. Players start out in $D1$, and when they are recorded playing D , they stochastically transition from $D1$ to P , then from P to S , and finally from S to $D2$, which is absorbing. Thus these strategies are a stochastic form of the strategy $DPSD$.

The intuition for why this works is that appropriately tuning the transition probabilities in stochastic $DPSD$ strategies is like tuning the parameters K and L in the $P_K S_L D_\infty$ strategies used to prove Theorem 3(b), but with greater flexibility because the tunable parameters are now real numbers rather than integers. This greater flexibility is what allows us to eliminate the nuisance term f . Furthermore, adding the first defector phase $D1$ at the beginning of the construction gives us additional flexibility to tune the population share of defectors, and in particular to keep it away from 0 when $\varepsilon \ll 1 - \gamma$. This is what lets us cover the double limit $(\gamma, \varepsilon) \rightarrow (1, 0)$ rather than only the iterated limit.¹⁷ Stochastic PSD strategies, which lack this first defector phase,

¹⁷Without PPR, adding a phase of defectors at the start of $P_K S_L D_\infty$ strategies enables efficiency in the *opposite* iterated limit, where first $\varepsilon \rightarrow 0$ and then $\gamma \rightarrow 1$. In the usual iterated limit, such strategies experience integer problems, so we do not know how much cooperation these strategies allow in the general double limit.

cannot achieve efficiency in the general $(\gamma, \varepsilon) \rightarrow (1, 0)$ limit, but can achieve efficiency in the usual iterated limit whenever $g < 1$ or $l > g + g^2$.

The $g < 1$ case of 4(b) is similar except that there are only two phases, a preciprocator phase P and a defector phase D , and players start out in P and stochastically transition to D , which is absorbing.

7 Characterization of Generalized Trigger Strategies

Our last result is that, with the Noisy Counting D 's record system, the $P_K S_L D_\infty$ generalized trigger strategies used to prove Theorem 3(b) are in fact the unique class of coordination-proof equilibria that satisfy the additional “forgery-proofness” condition that no player can gain by adding D 's to their record. This condition is realistic in settings where players can forge records of fake past interactions but cannot hide records of true interactions, such as an online rating system where users can post fake reviews to their profiles but cannot delete true reviews.¹⁸

Definition 8. *An equilibrium is **forgery-proof** if $V_k \geq V_{k'}$ for every $k \leq k'$.*

For example, cyclic strategies are not forgery-proof because any player with a defector record would forge D 's until reaching the next preciprocator record. On the other hand, generalized trigger strategies are forgery-proof because a player's flow payoff decreases as they progress from preciprocator records to supercooperator records and then to defector records, so forging D 's only speeds up the deterioration of a player's flow payoff. In fact, together with coordination-proofness, forgery-proofness characterizes generalized trigger strategies.

Theorem 5. *Assume the record system is Noisy Counting D 's. Any strict equilibrium that satisfies forgery-proofness and coordination-proofness corresponds to a generalized*

¹⁸In their lending model, Bhaskar and Thomas (2018) use a similar restriction to rule out equilibria where borrowers with two defaults are treated better than borrowers with only one.

trigger strategy. Moreover, if a generalized trigger strategy is a strict equilibrium, it satisfies forgery-proofness and coordination-proofness.

The result can be understood in three steps. First, forgery-proofness implies that there is a critical record \bar{k} such that all records $k \geq \bar{k}$ are defectors. Intuitively, this follows because, if there were infinitely many cooperator records, a player could profitably deviate by always playing D and then inflating their record to the next cooperator record. Second, this critical record \bar{k} can be chosen so that all records $k < \bar{k}$ are either preciprocators or supercooperators, because whenever there is a defector record that is followed by a cooperator record, a player at the defector record could profitably deviate by inflating their record to the next cooperator record. The third and last step is to classify the first \bar{k} records as preciprocators or supercooperators. Note that, whether a cooperator is a preciprocator or a supercooperator, their opponent will play the same way. Thus, all incentives to play C for players with records $k < \bar{k}$ come from avoiding the “punishment” of reaching record \bar{k} and triggering an increase in the fraction of partners who will play D . Since the survival probability γ is less than 1, this penalty looms larger the closer a player’s record is to \bar{k} . Hence, players with larger records are willing to incur greater costs to prevent their records from rising further. This implies that there is a critical record $k^* < \bar{k}$ such that only players with records greater than k^* are willing to play C at a cost of l (while players with records less than k^* are willing to play C at a cost of g , but not at a cost of l). We conclude that it is those players with lower records who must be preciprocators, which yields a $P_K S_L D_\infty$ profile with $K = k^*$ and $L = \bar{k} - k^*$.

In light of Theorem 5, it is interesting to note that trigger strategies of the form $P_K D_\infty$ (i.e., *GrimK*) can be equilibria when $g < \min\{l, 1\}$, where cyclic strategies can support limit efficiency but violate forgery-proofness. Moreover, whenever $g < l/(1+l)$, these trigger strategies can support a positive level of cooperation in the iterated limit where first $\gamma \rightarrow 1$ and then $\varepsilon \rightarrow 0$, and they can support limit efficiency if $l \rightarrow \infty$ (Clark, Fudenberg, and Wolitzky, 2019a), or if PPR is available for fixed l .

8 Discussion

This paper has analyzed the robust enforcement of cooperation in a large population with anonymous random matching and minimal information about partners' past play. Equilibria are required to be both strict and “coordination-proof,” in that societal cooperation cannot be based on threatened within-match miscoordination. We derived a sharp characterization: robust cooperation is possible when expected lifetimes are long and noise is small if and only if stage-game payoffs are either “mild” or “strongly supermodular.” The strength of the short-term coordination motive and the temptation to cheat thus determine the prospects for robust long-term cooperation.

We conclude with some observations about extensions and alternative models.

Non-coordination-proof equilibria. For some parameters, non-coordination-proof equilibria can support cooperation while coordination-proof equilibria cannot. Here is an example: The first four records ($k = 0, 1, 2, 3$) are reciprocator records, and all subsequent records ($k \geq 4$) are defector records. However, while the first two reciprocator records ($k = 0, 1$) play C against all reciprocators, the last two reciprocator records ($k = 3, 4$) play C only against those reciprocators with records $k = 0$ or $k = 1$. This strategy violates coordination-proofness. Moreover, for parameter values $\gamma = .892$ and $\varepsilon = .001$, there is an equilibrium with this strategy when $g = 1.01$ and $l = 1.95$. As we have seen, for these parameters cooperation cannot be supported by coordination-proof strategies. Characterizing when cooperation can be supported with general, non-coordination-proof strategies is an open question.

Higher-order information. This paper has restricted attention to first-order record systems. If records can also use second-order information, so that the update of a player's record can depend on their opponent's action as well as their own, supporting cooperation becomes much easier. Indeed, with second-order information CFW establish a Nash-threat folk theorem for general stage games without the need for any restriction on the stage-game parameters. Intuitively, second-order information can track whether a given play of D is an opportunistic deviation or a deserved punish-

ment, while first-order records cannot.

Less information. We have seen that, for almost any parameters for which any cooperation is possible for any first-order record system, limit efficiency is attainable for any record system that is richer than Noisy Counting D 's. Coarser record systems cannot always do as well: for example, if records count D 's only up to some fixed number then generalized trigger strategies can still be implemented but their efficiency is greatly reduced, and indeed goes to 0 as $\gamma \rightarrow 1$ for fixed ε .

Resetting records. A natural variant of Noisy Counting D 's arises when records are “lost” with some probability each period, so a player’s record is occasionally redrawn from some fixed distribution $\bar{\mu} \in \Delta(R)$. Our limit efficiency results can be shown to extend to the triple iterated limit where first the resetting probability goes to 0, then $\gamma \rightarrow 1$, and then $\varepsilon \rightarrow 0$.¹⁹ The situation is different if a player can reset their record at will, for example by re-entering the game under a pseudonym. In this case, society must use strategies where new players must “build a reputation” before anyone will cooperate with them, as in Friedman and Resnick (2001a).

Different stage games. In more general stage games it is harder to find necessary conditions for various equilibrium outcomes, but sufficient conditions can be obtained by construction. In this way CFW proves a Nash-threat folk theorem for general stage games with second-order records, and provides sufficient conditions for the enforceability of a given outcome using first-order records.

¹⁹In the special case where $\bar{\mu}$ is a degenerate distribution on the newborn-player record 0, introducing a resetting probability of p is equivalent to reducing the continuation probability to $\gamma(1-p)$, in which case the order of the limits $p \rightarrow 0$ and $\gamma \rightarrow 1$ makes no difference.

References

- [1] B Douglas Bernheim and Debraj Ray. “Collective Dynamic Consistency in Repeated Games”. *Games and Economic Behavior* (1989).
- [2] V Bhaskar and Caroline Thomas. “Community Enforcement of Trust with Bounded Memory”. *Review of Economic Studies* (2018).
- [3] Sylvain Chassang and Satoru Takahashi. “Robustness to incomplete information in repeated games”. *Theoretical Economics* (2011).
- [4] Daniel Clark, Drew Fudenberg, and Alexander Wolitzky. *Individualistic Records Robustly Support Social Cooperation*. Tech. rep. 2019.
- [5] Daniel Clark, Drew Fudenberg, and Alexander Wolitzky. *Steady-State Equilibria in Anonymous Repeated Games, I: Trigger Strategies in General Stage Games*. Tech. rep. 2019.
- [6] Pedro Dal Bó and Guillaume R Fréchette. “On the Determinants of Cooperation in Infinitely Repeated Games: A Survey”. *Journal of Economic Literature* (2018).
- [7] Joyee Deb, Takuo Sugaya, and Alexander Wolitzky. “The Folk Theorem in Repeated Games with Anonymous Random Matching”. 2018.
- [8] Francesc Dilmé. “Helping Behavior in Large Societies”. *International Economic Review* (2016).
- [9] Robert C Ellickson. *Order without law*. Harvard University Press, 1991.
- [10] Glenn Ellison. “Cooperation in the prisoner’s dilemma with anonymous random matching”. *Review of Economic Studies* (1994).
- [11] Joseph Farrell and Eric Maskin. “Renegotiation in Repeated Games”. *Games and Economic Behavior* (1989).
- [12] Eric J. Friedman and Paul Resnick. “The Social Cost of Cheap Pseudonyms”. *Journal of Economics and Management Strategy* (2001).
- [13] Eric J Friedman and Paul Resnick. “The social cost of cheap pseudonyms”. *Journal of Economics & Management Strategy* (2001).
- [14] Drew Fudenberg, David Levine, and Eric Maskin. “The Folk Theorem with Imperfect Public Information”. *Econometrica* (1994).
- [15] Drew Fudenberg, David G. Rand, and Anna Dreber. “Slow to Anger and Fast to Forgive: Cooperation in an Uncertain World”. *American Economic Review* (2012).
- [16] Avner Greif. “Contract enforceability and economic institutions in early trade: The Maghribi traders’ coalition”. *The American economic review* (1993).

- [17] Yuval Heller and Erik Mohlin. “Observations on Cooperation”. *Review of Economic Studies* (2018).
- [18] Johannes Hörner and Wojciech Olszewski. “The folk theorem for games with private almost-perfect monitoring”. *Econometrica* (2006).
- [19] Michihiro Kandori. “Social Norms and Community Enforcement”. *The Review of Economic Studies* (1992).
- [20] Benjamin Klein and Keith B Leffler. “The role of market forces in assuring contractual performance”. *Journal of political Economy* (1981).
- [21] Paul Milgrom, Douglass C. North, and Barry R. Weingast. “The Role of Institutions in the Revival of Trade: The Law Merchant, Private Judges, and the Champagne Fairs”. *Economics and Politics* (1990).
- [22] Martin A. Nowak and Karl Sigmund. “Evolution of indirect reciprocity by image scoring”. *Nature* (1998).
- [23] Elinor Ostrom. *Governing the commons: The evolution of institutions for collective action*. Cambridge university press, 1990.
- [24] Karthik Panchanathan and Robert Boyd. “A tale of two defectors: the importance of standing for evolution of indirect reciprocity”. *Journal of Theoretical Biology* (2003).
- [25] Robert W. Rosenthal. “Sequences of Games with Varying Opponents”. *Econometrica* (1979).
- [26] Robert W. Rosenthal and Henry J. Landau. “A game-theoretic analysis of bargaining with reputations”. *Journal of Mathematical Psychology* (1979).
- [27] Satoru Takahashi. “Community enforcement when players observe partners’ past play”. *Journal of Economic Theory* (2010).

Appendix

All omitted proofs are in the Online Appendix (OA).

A.1 Proof of Lemma 6

Let $\bar{V} = \sup_r V_r$ and let $\{r_n\}_{n \in \mathbb{N}}$ be a sequence of records such that $\lim_{n \rightarrow \infty} V_{r_n} = \bar{V}$. Note that $\bar{V} < \infty$ and, since V_0 (the expected lifetime payoff of a newborn player) equals $\mu^P \mu^C + \mu^S(\mu^C - \mu^D l) + \mu^D \mu^S(1 + g)$ (the average flow payoff in the population),

we have $\bar{V} \geq V_0 = \mu^P \mu^C + \mu^S(\mu^C - \mu^D l) + \mu^D \mu^S(1 + g)$.

Lemma 11. *If $\mu^C > 0$, there is no sequence of defector records $\{r_n\}_{n \in \mathbb{N}}$ such that $\lim_{n \rightarrow \infty} V_{r_n} = \bar{V}$.*

Proof. Suppose that there is a sequence of defector records $\{r_n\}_{n \in \mathbb{N}}$ such that $\lim_{n \rightarrow \infty} V_{r_n} = \bar{V}$. Since $V_r = (1 - \gamma)\pi_r + \gamma(1 - p_r^D)V_r^C + \gamma p_r^D V_r^D$ and $\pi_{r_n} = \mu^S(1 + g)$ for all r_n , we have $V_{r_n} = (1 - \gamma)\mu^S(1 + g) + \gamma(1 - p_{r_n}^D)V_{r_n}^C + \gamma p_{r_n}^D V_{r_n}^D$ for all r_n . This implies

$$V_{r_n} \leq \mu^S(1 + g) + \frac{\gamma}{1 - \gamma}(1 - p_{r_n}^D) \max\{V_{r_n}^C - V_{r_n}, 0\} + \frac{\gamma}{1 - \gamma} p_{r_n}^D \max\{V_{r_n}^D - V_{r_n}, 0\}.$$

Since $\lim_{n \rightarrow \infty} V_{r_n} = \bar{V}$, $\lim_{n \rightarrow \infty} \max\{V_{r_n}^C - V_{r_n}, 0\} = \lim_{n \rightarrow \infty} \max\{V_{r_n}^D - V_{r_n}, 0\} = 0$. It further follows that $\bar{V} = \lim_{n \rightarrow \infty} V_{r_n} \leq \mu^S(1 + g)$, so $V_r \leq \mu^S(1 + g)$ for all r . However, note that every player can secure an expected flow payoff of $\mu^S(1 + g)$ every period by always defecting, so it must be that $V_r \geq \mu^S(1 + g)$ for all r . It follows that $V_r = \mu^S(1 + g)$ for all r , and since the value function is constant across records, every record must be a defector record, so $\mu^C = 0$. \blacksquare

Lemma 12. *If $\mu^C > 0$, there is some record r' that is a preciprocator or a supercooperator and is such that*

$$V_{r'} - \frac{\gamma}{1 - \gamma}(V_{r'}^C - V_{r'}) \geq \mu^P \mu^C + \mu^S(\mu^C - \mu^D l) + \mu^D \mu^S(1 + g). \quad (2)$$

Proof. First, consider the case where $\bar{V} = \mu^P \mu^C + \mu^S(\mu^C - \mu^D l) + \mu^D \mu^S(1 + g)$. Then there must be some record r' such that $V_{r'} = \mu^P \mu^C + \mu^S(\mu^C - \mu^D l) + \mu^D \mu^S(1 + g)$. By Lemma 11, such a r' cannot be a defector record and so must be either a preciprocator or a supercooperator. Additionally, $V_{r'}^C \leq \bar{V}$, so $V_{r'} - (\gamma/(1 - \gamma))(V_{r'}^C - V_{r'}) \geq \mu^P \mu^C + \mu^S(\mu^C - \mu^D l) + \mu^D \mu^S(1 + g)$.

Now, consider the case where $\bar{V} > \mu^P \mu^C + \mu^S(\mu^C - \mu^D l) + \mu^D \mu^S(1 + g)$. For any sequence of records $\{r_n\}_{n \in \mathbb{N}}$ such that $\lim_{n \rightarrow \infty} V_{r_n} = \bar{V}$, $\lim_{n \rightarrow \infty} \max\{V_{r_n}^C - V_{r_n}, 0\} = 0$, so there is some sufficiently high n such that $V_{r_n} - (\gamma/(1 - \gamma))(V_{r_n}^C - V_{r_n}) \geq \mu^P \mu^C +$

$\mu^S(\mu^C - \mu^D l) + \mu^D \mu^S(1 + g)$. Additionally, by Lemma 11, for sufficiently high n , the record r_n must be either a reciprocator or a supercooperator. ■

We now present the proof of Lemma 6.

Lemma 6. *In any coordination-proof equilibrium with $\mu^C > 0$, $\mu^P + \mu^S(l - g) > g$.*

Proof. First, take the case where r' is a reciprocator. Then by Lemma 3, we must have

$$\frac{1 - \nu_C(r') - \nu_D(r')}{p_{r'}^D} \left(\pi_{r'} - V_{r'} + \frac{\gamma}{1 - \gamma} (V_{r'}^C - V_{r'}) \right) > g.$$

When $\pi_{r'} = \mu^C$ and $V_{r'} - \frac{\gamma}{1 - \gamma} (V_{r'}^C - V_{r'}) \geq \mu^P \mu^C + \mu^S(\mu^C - \mu^D l) + \mu^D \mu^S(1 + g)$, this implies

$$\frac{(1 - \nu_C(r') - \nu_D(r')) \mu^D}{p_{r'}^D} (\mu^P + \mu^S(l - g)) > g.$$

Note that $p_{r'}^D \geq (1 - \nu_D(r')) \mu^D$ since a reciprocator plays D whenever they are matched with a defector and this leads to a recording of D being fed into the record system with probability $1 - \nu_D(r')$. This gives $(1 - \nu_C(r') - \nu_D(r')) \mu^D / p_{r'}^D < 1$, so $\mu^P + \mu^S(l - g) > g$ must hold.

Now, take the case where r' is a supercooperator. Then by Lemma 3, we must have $\pi_{r'} - V_{r'} + (\gamma / (1 - \gamma))(V_{r'}^C - V_{r'}) > 0$. When $\pi_{r'} = \mu^C - \mu^D l$ and $V_{r'} - \frac{\gamma}{1 - \gamma} (V_{r'}^C - V_{r'}) \geq \mu^P \mu^C + \mu^S(\mu^C - \mu^D l) + \mu^D \mu^S(1 + g)$, this implies

$$\mu^C - \mu^D l - (\mu^P \mu^C + \mu^S(\mu^C - \mu^D l) + \mu^D \mu^S(1 + g)) = \mu^D (\mu^P + \mu^S(l - g) - l) > 0.$$

This requires $\mu^P + \mu^S(l - g) > l$, which implies $\mu^P + \mu^S(l - g) > g$, since $l > g$. ■

A.2 Proofs of Lemmas 7 and 8

Lemma 7. *With cyclic strategies, let $\mu^C(\gamma, \varepsilon)$ denote the share of cooperators in some steady state for arbitrary $(\gamma, \varepsilon) \in (0, 1) \times (0, 1)$. Then $\lim_{\gamma \rightarrow 1} \mu^C(\gamma, \varepsilon) = \bar{\mu}^C(\varepsilon)$.*

Proof of Lemma 7. Let $\mu_i(\gamma, \varepsilon)$ be the share of players in phase i for some steady state at parameters (γ, ε) , and similarly let $\mu^C(\gamma, \varepsilon) = \sum_{i=0}^{K-1} \mu_i(\gamma, \varepsilon)$ be the corresponding steady-state share of cooperators. Now fix ε , and let $\{\gamma_j\}_{j \in \mathbb{N}} \subseteq \mathbb{N}$ be some sequence of γ in which $\lim_{j \rightarrow \infty} \gamma_j = 1$ and each corresponding sequence $\mu_i(\gamma_j, \varepsilon)$ converges to some $\tilde{\mu}_i(\varepsilon)$. Let $\tilde{\mu}^C(\varepsilon) = \sum_{i=0}^{K-1} \tilde{\mu}_i(\varepsilon)$ denote the corresponding limit of the share of cooperators. We will show that $\tilde{\mu}^C(\varepsilon) = \bar{\mu}^C(\varepsilon)$. Since this holds for any convergent subsequence, we conclude that $\lim_{\gamma \rightarrow 1} \mu^C(\gamma, \varepsilon) = \bar{\mu}^C(\varepsilon)$.

Note that the total outflow from the preciprocator phases is the share of preciprocators that die in a given round plus the share of phase $K - 1$ preciprocators that transition to phase K , which is $(1 - \gamma_j)\mu^C(\gamma_j, \varepsilon) + \gamma_j(1 - (1 - \varepsilon)\mu^C(\gamma_j, \varepsilon))\mu_{K-1}(\gamma_j, \varepsilon)$. Moreover, the total inflow into the preciprocator phases is the share of newborn players plus the share of phase $K + M - 1$ defectors that transition back to phase 0. Formally, this is $1 - \gamma_j + \gamma_j\mu_{K+M-1}(\gamma_j, \varepsilon)$. Setting these expressions equal to each other and taking the limit as $j \rightarrow \infty$ yields

$$(1 - (1 - \varepsilon)\tilde{\mu}^C(\varepsilon))\tilde{\mu}_{K-1}(\varepsilon) = \tilde{\mu}_{K+M-1}(\varepsilon). \quad (3)$$

Note that when $0 < i \leq K - 1$, both phase i and phase $i - 1$ are preciprocators. Therefore, the outflow from phase i is $(1 - \gamma(1 - \varepsilon)\mu^C(\gamma, \varepsilon))\mu_i(\gamma, \varepsilon)$, while the inflow into phase i is $\gamma(1 - (1 - \varepsilon)\mu^C(\gamma, \varepsilon))\mu_{i-1}(\gamma, \varepsilon)$. Thus,

$$\mu_i(\gamma_j, \varepsilon) = \frac{\gamma_j(1 - (1 - \varepsilon)\mu^C(\gamma_j, \varepsilon))}{1 - \gamma(1 - \varepsilon)\mu^C(\gamma_j, \varepsilon)}\mu_{i-1}(\gamma, \varepsilon) = \beta(\gamma_j, \varepsilon, \mu^C)\mu_{i-1}(\gamma_j, \varepsilon),$$

which gives $\tilde{\mu}_i(\varepsilon) = \tilde{\mu}_{i-1}(\varepsilon)$, since $\lim_{j \rightarrow \infty} \beta(\gamma_j, \varepsilon, \mu^C(\gamma_j, \varepsilon)) = 1$. Since this holds for all $0 < i \leq K - 1$, we conclude that $\tilde{\mu}_i(\varepsilon) = \tilde{\mu}_{K-1}(\varepsilon)$ for all $0 \leq i \leq K - 1$, so

$$\tilde{\mu}_{K-1}(\varepsilon) = \frac{1}{K}\tilde{\mu}^C(\varepsilon). \quad (4)$$

When $K < i \leq K + M - 1$, both phase i and phase $i - 1$ are defectors. Therefore,

the outflow from phase i is $\mu_i(\gamma, \varepsilon)$, while the inflow into phase i is $\gamma\mu_{i-1}(\gamma, \varepsilon)$. Thus, $\mu_i(\gamma_j, \varepsilon) = \gamma_j\mu_{i-1}(\gamma, \varepsilon)$, which consequently gives $\tilde{\mu}_i(\varepsilon) = \tilde{\mu}_{i-1}(\varepsilon)$. Since this holds for all $K < i \leq K+M-1$, we conclude that $\tilde{\mu}_i(\varepsilon) = \tilde{\mu}_{K+M-1}(\varepsilon)$ for all $K \leq i \leq K+M-1$, so

$$\tilde{\mu}_{K+M-1}(\varepsilon) = \frac{1}{M}(1 - \tilde{\mu}^C(\varepsilon)). \quad (5)$$

By Equations 3, 4, and 5, we obtain

$$\frac{1}{K}(1 - (1 - \varepsilon)\tilde{\mu}^C(\varepsilon))\tilde{\mu}^C(\varepsilon) = \frac{1}{M}(1 - \tilde{\mu}^C(\varepsilon)),$$

and solving this for $\tilde{\mu}^C(\varepsilon)$ gives $\tilde{\mu}^C(\varepsilon) = \bar{\mu}^C(\varepsilon)$. ■

Lemma 8. *With cyclic strategies, let $\mu^C(\gamma, \varepsilon)$ denote the share of cooperators and let $V_i(\gamma, \varepsilon)$ denote the value function for phase i in some steady state for arbitrary $(\gamma, \varepsilon) \in (0, 1) \times (0, 1)$. Then,*

$$\lim_{\varepsilon \rightarrow 0} \lim_{\gamma \rightarrow 1} \frac{1 - \varepsilon}{1 - (1 - \varepsilon)\mu^C(\gamma, \varepsilon)} (\mu^C(\gamma, \varepsilon) - V_i(\gamma, \varepsilon)) = \rho$$

for all $0 \leq i \leq K - 1$.

Proof of Lemma 8. First, note that for all $0 \leq i < K - 1$,

$$V_i(\gamma, \varepsilon) = (1 - \gamma)\mu^C(\gamma, \varepsilon) + \gamma(1 - \varepsilon)\mu^C(\gamma, \varepsilon)V_i(\gamma, \varepsilon) + \gamma(1 - (1 - \varepsilon)\mu^C(\gamma, \varepsilon))V_{i+1}(\gamma, \varepsilon).$$

Taking the limit as $\gamma \rightarrow 1$ gives $\lim_{\gamma \rightarrow 1} V_i(\gamma, \varepsilon) = \lim_{\gamma \rightarrow 1} V_{i+1}(\gamma, \varepsilon)$, should the limits exist. Since this holds for all $0 \leq i < K - 1$ and, by Lemma 7, $V_0(\gamma, \varepsilon) = (\mu^C(\gamma, \varepsilon))^2$ implies that $\lim_{\gamma \rightarrow 1} V_0(\gamma, \varepsilon) = (\bar{\mu}^C(\varepsilon))^2$, we conclude that $\lim_{\gamma \rightarrow 1} V_i(\gamma, \varepsilon) = (\bar{\mu}^C(\varepsilon))^2$ for all $0 \leq i \leq K - 1$.

Thus, all that remains to be shown is that

$$\lim_{\varepsilon \rightarrow 0} \frac{1 - \varepsilon}{1 - (1 - \varepsilon)\bar{\mu}^C(\varepsilon)} (\bar{\mu}^C(\varepsilon) - (\bar{\mu}^C(\varepsilon))^2) = \rho. \quad (6)$$

Note that

$$\frac{1 - \varepsilon}{1 - (1 - \varepsilon)\bar{\mu}^C(\varepsilon)}(\bar{\mu}^C(\varepsilon) - (\bar{\mu}^C(\varepsilon))^2) = (1 - \varepsilon) \left(\frac{\frac{1 - \bar{\mu}^C(\varepsilon)}{\varepsilon}}{\bar{\mu}^C(\varepsilon) + \frac{1 - \bar{\mu}^C(\varepsilon)}{\varepsilon}} \right) \bar{\mu}^C(\varepsilon).$$

Since $\lim_{\varepsilon \rightarrow 0} \bar{\mu}^C(\varepsilon) = 1$ and $\lim_{\varepsilon \rightarrow 0} (1 - \bar{\mu}^C(\varepsilon))/\varepsilon = \rho/(1 - \rho)$, as can be readily confirmed, we conclude that Equation 6 holds. \blacksquare

A.3 Proof of Lemma 10

We define two functions, $\alpha : (0, 1) \times (0, 1) \rightarrow (0, 1)$ and $\beta : (0, 1) \times (0, 1) \times [0, 1] \rightarrow (0, 1)$, where

$$\begin{aligned} \alpha(\gamma, \varepsilon) &= \frac{\gamma\varepsilon}{1 - \gamma(1 - \varepsilon)}, \\ \beta(\gamma, \varepsilon, \mu) &= \frac{\gamma(1 - (1 - \varepsilon)\mu)}{1 - \gamma(1 - \varepsilon)\mu}. \end{aligned} \tag{7}$$

Lemma 13. *There is a $P_K S_L D_\infty$ equilibrium with share of cooperators μ^C , share of preciprocators μ^P , and share of supercooperators μ^S if and only if the following conditions hold:*

1. *Feasibility:*

$$\begin{aligned} \mu^C &= 1 - \alpha(\gamma, \varepsilon)^L \beta(\gamma, \varepsilon, \mu^C)^K, \\ \mu^P &= 1 - \beta(\gamma, \varepsilon, \mu^C)^K, \\ \mu^S &= (1 - \alpha(\gamma, \varepsilon)^L) \beta(\gamma, \varepsilon, \mu^C)^K. \end{aligned}$$

2. *Incentives:*

$$\begin{aligned} (C|C)_0 &: \frac{(1 - \varepsilon)(1 - \mu^C)}{1 - (1 - \varepsilon)\mu^C} (\mu^P + \mu^S(l - g)) > g, \\ (D|D)_{K-1} &: \frac{\gamma(1 - \varepsilon)(1 - \mu^C)}{1 - \gamma(1 - \varepsilon)\mu^C} (\mu^P + \mu^S(l - g)) + \mu^P l < l, \\ (C|D)_K \text{ (if } \mu^S > 0) &: \frac{(1 - \varepsilon)(1 - \mu^C)}{1 - (1 - \varepsilon)\mu^C} (\mu^P + \mu^S(l - g)) + \mu^P l > l. \end{aligned}$$

The proof of Lemma 13 is in OA.1. The feasibility constraints come from calculating the steady-state shares μ_k for the generalized trigger strategy $P_K S_L D_\infty$ as a function of μ^C and then setting $\mu^C = \sum_{k=0}^{K+L-1} \mu_k$, $\mu^P = \sum_{k=0}^{K-1} \mu_k$, and $\mu^S = \sum_{k=K}^{K+L-1} \mu_k$. The $(C|C)_0$ and $(C|D)_K$ incentive constraints come from solving V_0 and V_K , the value functions at the corresponding records, and using Lemma 4, while the $(D|D)_{K-1}$ constraint is derived by using the value of V_K , after relating V_{K-1} and V_K .

Fix $\mu^P \in (g/(1+g), 1 - g/l]$. Consider the equation

$$\frac{(1-\varepsilon)(1-\mu^C)}{1-(1-\varepsilon)\mu^C} ((l-g)\mu^C + (1+g-l)\mu^P) + l\mu^P = l \quad (8)$$

and the function $h(\varepsilon, \mu^P)$ defined by $h(\varepsilon, \mu^P) = \max\{\mu^C \in [0, 1] : \mu^C \text{ solves Equation 8}\}$. If $h(\varepsilon, \mu^P)$ is well-defined, it gives the maximum level of cooperation for the given μ^P and ε that satisfies the $\gamma \rightarrow 1$ “limit” incentive constraints of Lemma 13. Straightforward calculations show that, for any $\mu^P \in (g/(1+g), 1 - g/l]$, $h(\varepsilon, \mu^P)$ is well-defined for sufficiently small and positive ε , and that

$$\lim_{\varepsilon \rightarrow 0} \frac{1 - h(\varepsilon, \mu^P)}{\varepsilon} = \frac{l(1 - \mu^P)}{(1+g)\mu^P - g}.$$

An immediate implication of this is $\lim_{\varepsilon \rightarrow 0} h(\varepsilon, \mu^P) = 1$. Combining this with the following two lemmas proves Lemma 10.

Let $\kappa : (g/(1+g), 1 - g/l] \rightarrow \mathbb{R}$ be the function given by

$$\kappa(\mu^P) = \frac{l \ln(1 - \mu^P)(1 - \mu^P)}{l - g + (1 + g - l)\mu^P},$$

and $\iota : (g/(1+g), 1 - g/l] \rightarrow \mathbb{R}_+$ be the function given by

$$\iota(\mu^P) = \frac{(1+g)\mu^P - g + 1}{l - g + (1 + g - l)\mu^P}.$$

Let $\bar{\mu}_{trig}^C(\gamma, \varepsilon)$ be the maximal share of cooperators in any equilibrium using any

$P_K S_L D_\infty$ generalized trigger strategy for parameters γ and ε :

$$\bar{\mu}_{trig}^C(\gamma, \varepsilon) = \sup \{ \mu^C : \mu^C \text{ is the share of cooperators in a } P_K S_L D_\infty \text{ equilibrium} \}.$$

Lemma 14. *Fix $\mu^P \in (g/(1+g), 1 - g/l]$. If $|1 + \kappa(\mu^P)| > \iota(\mu^P)$, then there exists some $\bar{\varepsilon} > 0$ such that $\liminf_{\gamma \rightarrow 1} \bar{\mu}_{trig}^C(\gamma, \varepsilon) \geq h(\varepsilon, \mu^P)$ for $\varepsilon < \bar{\varepsilon}$.*

Lemma 15. *Suppose that $l > g + g^2$. Some $\mu^P \in (g/(1+g), 1 - g/l]$ satisfies $|1 + \kappa(\mu^P)| > \iota(\mu^P)$ if $l > \max\{g + g^2, f(g)\}$.*

OA.2.1 presents the proof of Lemma 14. It uses the inverse function theorem and other tools of differential calculus to show that, when $|1 + \kappa(\mu^P)| > \iota(\mu^P)$, for sufficiently small ε , any neighborhood of $(h(\varepsilon, \mu^P), \mu^P)$ can be approached by feasible profiles for sufficiently high γ . The proof of Lemma 15 is in OA.2.2.

A.4 Limit Efficiency Using PPR's

A.4.1 Limit Efficiency for $g < 1$

We use the class of stochastic $P_{\chi_P} D$ strategies, where players start out in the pre-ciprocator phase P , and, when they are recorded as playing D , transition to D with probability χ_P . The defector phase D is absorbing. We first characterize the population shares possible in equilibria with this class of strategies.

Lemma 16. *There is a $P_{\chi_P} D$ equilibrium with share μ^C of players in P iff the following feasibility constraint*

$$\chi_P = \frac{(1 - \gamma)(1 - \mu^C)}{\gamma(1 - (1 - \varepsilon)\mu^C)\mu^C} \leq 1$$

and incentive constraint

$$g < \frac{(1 - \varepsilon)(1 - \mu^C)}{1 - (1 - \varepsilon)\mu^C} \mu^C < l$$

are satisfied.

The proof of Lemma 16 is in OA.3. The feasibility constraint comes from calculating the transition probability χ_P necessary to support a steady-state share of μ^C in P . The incentive constraint comes from solving V^P , the value function of players in P , and using Lemma 4.

We now show that we can attain efficiency in the general $(\gamma, \varepsilon) \rightarrow (1, 0)$ limit whenever $g < 1$.

Lemma 17. *If $g < 1$, then there is a sequence of equilibria in which $\lim_{(\gamma, \varepsilon) \rightarrow (1, 0)} \mu^C(\gamma, \varepsilon) = 1$.*

Proof of Lemma 17. Fix some κ such that $g < \kappa < \min\{l, 1\}$. Consider supporting $\mu^C(\gamma, \varepsilon) = \tilde{\mu}^C(\varepsilon)$, defined by

$$\tilde{\mu}^C(\varepsilon) = \frac{(1 - \varepsilon)(1 + \kappa) + \sqrt{(1 - \varepsilon)^2(1 + \kappa)^2 - 4\kappa(1 - \varepsilon)}}{2(1 - \varepsilon)\kappa}.$$

$\tilde{\mu}^C(\varepsilon)$ is well-defined for sufficiently small ε and satisfies

$$\frac{(1 - \varepsilon)(1 - \tilde{\mu}^C(\varepsilon))}{1 - (1 - \varepsilon)\tilde{\mu}^C(\varepsilon)} \tilde{\mu}^C(\varepsilon) = \kappa.$$

Additionally, $\lim_{\varepsilon \rightarrow 0} \tilde{\mu}^C(\varepsilon) = 1$, so the lemma follows if

$$\frac{(1 - \gamma)(1 - \tilde{\mu}^C(\varepsilon))}{\gamma(1 - (1 - \varepsilon)\tilde{\mu}^C(\varepsilon))\tilde{\mu}^C(\varepsilon)} \leq 1$$

for all (γ, ε) sufficient close to $(1, 0)$. That this is the case follows from combining

$$\frac{1 - \tilde{\mu}^C(\varepsilon)}{1 - (1 - \varepsilon)\tilde{\mu}^C(\varepsilon)} = \frac{1 - \tilde{\mu}^C(\varepsilon)}{1 - \tilde{\mu}^C(\varepsilon)^C + \varepsilon\tilde{\mu}^C(\varepsilon)} \leq 1$$

for all (γ, ε) with the fact that $\lim_{(\gamma, \varepsilon) \rightarrow (1, 0)} (1 - \gamma)/(\gamma\tilde{\mu}^C(\varepsilon)) = 0$. ■

A.4.2 Limit Efficiency for $l > g + g^2$

We use the class of stochastic $D_{\chi_{D_1}}P_{\chi_P}S_{\chi_S}D$ strategies, where players start out in the defector phase D_1 , and, when they are recorded as playing D , transition to the reciprocator phase P with probability χ_{D_1} . When a player in P is recorded as playing D , they transition to the supercooperator phase S with probability χ_P . Finally, when a player in S is recorded as playing D , they transition to the defector phase D_2 with probability χ_S . The phase D_2 is absorbing. We first characterize the population shares possible in equilibria with this class of strategies.

Lemma 18. *There is a $D_{\chi_{D_1}}P_{\chi_P}S_{\chi_S}D$ equilibrium with share μ^{D_1} of players in D_1 , share μ^P of players in P , share μ^S of players in S , and share μ^{D_2} of players in D_2 iff the following feasibility constraints*

$$\begin{aligned}\chi_{D_1} &= \frac{(1-\gamma)(1-\mu^{D_1})}{\gamma\mu^{D_1}} \leq 1, \\ \chi_P &= \frac{(1-\gamma)(1-\mu^{D_1}-\mu^P)}{\gamma(1-(1-\varepsilon)\mu^C)\mu^P} \leq 1, \\ \chi_S &= \frac{(1-\gamma)\mu^{D_2}}{\gamma\varepsilon\mu^S} \leq 1,\end{aligned}$$

and incentive constraints

$$\begin{aligned}P : g &< \frac{1}{1-\mu^{D_1}} \frac{(1-\varepsilon)(1-\mu^C)}{1-(1-\varepsilon)\mu^C} \left(\mu^C + \mu^S l - \frac{\mu^{D_1}}{\mu^D} \mu^C - \frac{\mu^{D_2}}{\mu^D} \mu^S (1+g) \right) < l, \\ S : \frac{1-\varepsilon}{\varepsilon} \frac{\mu^{D_2}}{\mu^S + \mu^{D_2}} (\mu^C - \mu^D l - \mu^S (1+g)) &> l,\end{aligned}$$

are satisfied.

The proof of Lemma 18 is in OA.4. The feasibility constraints come from calculating the transition probabilities χ_{D_1} , χ_P , and χ_S necessary to support the steady-state share of μ^{D_1} in D_1 , μ^P in P , and μ^S in S . The incentive constraints come from solving V^P and V^S , the value function of players in P and S , respectively, and then using Lemma 4.

Consider the following inequalities:

$$\begin{aligned}\mu^C + \mu^S(l - g - 1) &> g, \\ \mu^C - \mu^D l - \mu^S(1 + g) &> 0.\end{aligned}\tag{9}$$

Let $A(g, l)$ be the set of $\mu^C \in (0, 1)$ such that there exists some $g/l < \mu^S < \mu^C$ where the above inequalities are satisfied.

Lemma 19. *When $l > g + g^2$, $\sup\{A(g, l)\} = 1$.*

Proof of Lemma 19. Take $\mu^C = 1 - \kappa_1$ and $\mu^S = 1/(1 + g) - \kappa_2$ for small $\kappa_1, \kappa_2 > 0$. Note that $g/l < \mu^S < \mu^C$ for sufficiently small κ_1 and κ_2 . As $\kappa_1 \rightarrow 0$, $\mu^C \rightarrow 1$. Also, for sufficiently small $\kappa_2 > 0$,

$$\begin{aligned}\mu^C + \mu^S(l - g - 1) &\rightarrow \frac{l}{1 + g} - \kappa_2(l - g - 1) > g, \\ \mu^C - \mu^D l - \mu^S(1 + g) &\rightarrow \kappa_2(1 + g) > 0,\end{aligned}$$

where the first inequality follows directly from $l > g + g^2$. ■

We now show that we can attain efficiency in the general $(\gamma, \varepsilon) \rightarrow (1, 0)$ limit whenever $l > g + g^2$.

Lemma 20. *If $l > g + g^2$, there is a sequence of equilibria in which $\lim_{(\gamma, \varepsilon) \rightarrow (1, 0)} \mu^C(\gamma, \varepsilon) = 1$.*

Proof of Lemma 20. Fix some $\mu^C \in A(g, l)$ and some corresponding $g/l < \mu^S < \mu^C$ such that the inequalities in (9) are satisfied. Consider supporting such μ^C and μ^S along with

$$\begin{aligned}\mu^{D_2}(\gamma, \varepsilon) &= \lambda \min \left\{ 1, \frac{\gamma}{1 - \gamma} \varepsilon \right\} \mu^S, \\ \mu^{D_1}(\gamma, \varepsilon) &= \mu^C - \mu^{D_2}(\gamma, \varepsilon),\end{aligned}$$

where $0 < \lambda < 1$ is taken to be sufficiently small so that $\mu^P + (1 + \lambda)\mu^S < 1$. We will argue that for all (γ, ε) sufficiently close to $(1, 0)$, such shares can be supported in equilibrium. The conclusion then follows since, by Lemma 19, there are $\mu^C \in A(g, l)$ that are arbitrarily close to 1.

First, we show that the feasibility constraints are satisfied for (γ, ε) sufficiently close to $(1, 0)$. Note that $\mu^{D_1}(\gamma, \varepsilon) \geq 1 - \mu^P - (1 + \lambda)\mu^S$ regardless of (γ, ε) . Thus,

$$\lim_{(\gamma, \varepsilon) \rightarrow (1, 0)} \frac{(1 - \gamma)(1 - \mu^{D_1}(\gamma, \varepsilon))}{\gamma \mu^{D_1}(\gamma, \varepsilon)} = 0.$$

Additionally,

$$\lim_{(\gamma, \varepsilon) \rightarrow (1, 0)} \frac{(1 - \gamma)(1 - \mu^{D_1}(\gamma, \varepsilon) - \mu^P)}{\gamma(1 - (1 - \varepsilon)\mu^C)\mu^P} = 0,$$

and

$$\frac{(1 - \gamma)\mu^{D_2}(\gamma, \varepsilon)}{\gamma \varepsilon \mu^S} = \lambda \min \left\{ 1, \frac{\gamma}{1 - \gamma} \varepsilon \right\} \leq \lambda$$

for all (γ, ε) .

Now we show that the incentive constraints are satisfied for (γ, ε) sufficiently close to $(1, 0)$. Note that

$$\mu^C + \mu^S l - \frac{\mu^{D_1}(\gamma, \varepsilon)}{\mu^D} \mu^C - \frac{\mu^{D_2}(\gamma, \varepsilon)}{\mu^D} \mu^S (1 + g) \geq \min \{ \mu^S l, \mu^C + \mu^S (l - g - 1) \}.$$

Since $\mu^S > g/l$, $\mu^C + \mu^S (l - g - 1) > g$, and $\lim_{\varepsilon \rightarrow 0} (1 - \varepsilon)(1 - \mu^C)/(1 - (1 - \varepsilon)\mu^C) = 1$, it follows that

$$\left(\frac{1}{1 - \mu^{D_1}(\gamma, \varepsilon)} \right) \frac{(1 - \varepsilon)(1 - \mu^C)}{1 - (1 - \varepsilon)\mu^C} \left(\mu^C + \mu^S l - \frac{\mu^{D_1}(\gamma, \varepsilon)}{\mu^D} \mu^C - \frac{\mu^{D_2}(\gamma, \varepsilon)}{\mu^D} \mu^S (1 + g) \right) > g$$

for sufficiently small ε . Moreover,

$$\mu^C + \mu^S l - \frac{\mu^{D_1}(\gamma, \varepsilon)}{\mu^D} \mu^C - \frac{\mu^{D_2}(\gamma, \varepsilon)}{\mu^D} \mu^S (1 + g) \leq \max \{ \mu^S l, \mu^C + \mu^S (l - g - 1) \}.$$

Since $(1 - \varepsilon)(1 - \mu^C)/(1 - (1 - \varepsilon)\mu^C) < 1$ and $\mu^S < \mu^C < 1 - \mu^{D_1}(\gamma, \varepsilon)$, it follows that

$$\left(\frac{1}{1 - \mu^{D_1}(\gamma, \varepsilon)} \right) \frac{(1 - \varepsilon)(1 - \mu^C)}{1 - (1 - \varepsilon)\mu^C} \left(\mu^C + \mu^S l - \frac{\mu^{D_1}(\gamma, \varepsilon)}{\mu^D} \mu^C - \frac{\mu^{D_2}(\gamma, \varepsilon)}{\mu^D} \mu^S (1 + g) \right) < l$$

for all (γ, ε) .

Furthermore,

$$\frac{1 - \varepsilon}{\varepsilon} \frac{\mu^{D_2}(\gamma, \varepsilon)}{\mu^S + \mu^{D_2}(\gamma, \varepsilon)} \geq \frac{\lambda}{1 + \lambda} \min \left\{ \frac{1 - \varepsilon}{\varepsilon}, (1 - \varepsilon) \frac{\gamma}{1 - \gamma} \right\}.$$

Since $\lim_{\varepsilon \rightarrow 0} (1 - \varepsilon)/\varepsilon = \infty$ and $\lim_{(\gamma, \varepsilon) \rightarrow (1, 0)} (1 - \varepsilon)\gamma/(1 - \gamma) = \infty$, it follows that

$$\frac{1 - \varepsilon}{\varepsilon} \frac{\mu^{D_2}(\gamma, \varepsilon)}{\mu^S + \mu^{D_2}(\gamma, \varepsilon)} (\mu^C - \mu^D l - \mu^S (1 + g)) > l$$

for all (γ, ε) sufficiently close to $(1, 0)$ since $\mu^C - \mu^D l - \mu^S (1 + g) > 0$. ■

A.5 Proof of Theorem 5

The proof that any equilibrium that satisfies forgery-proofness and coordination-proofness corresponds to a generalized trigger strategy proceeds by establishing the following two lemmas.

Lemma 21. *In any equilibrium that satisfies forgery-proofness, if record k' is a defector, then any record $k \geq k'$ is also a defector.*

Proof of Lemma 21. Suppose that record k' is a defector. Then

$$V_{k'} = (1 - \gamma)\mu^S(1 + g) + \gamma V_{k'+1}.$$

By forgery-proofness, $V_{k'+1} \leq V_{k'}$, so it must be that $V_{k'} \leq \mu^S(1 + g)$, and consequently $V_k \leq \mu^S(1 + g)$ for all $k \geq k'$. Since every player can secure an expected flow payoff of $\mu^S(1 + g)$ in every period by always playing D , the value function at every record

must weakly exceed $\mu^S(1 + g)$, so we conclude that $V_k = \mu^S(1 + g)$ for every $k \geq k'$. Since the value function is constant for all records $k \geq k'$, every record $k \geq k'$ is a defector. \blacksquare

Lemma 22. *In any equilibrium, if record k is a supercooperator, then record $k + 1$ is not a preciprocator.*

Proof of Lemma 22. Suppose towards a contradiction that there exists k such that record k is a supercooperator and record $k + 1$ is a preciprocator. Then

$$V_k = (1 - \gamma)(\mu^C - \mu^D l) + \gamma(1 - \varepsilon)V_k + \gamma\varepsilon V_{k+1},$$

so $V_k = (1 - \alpha(\gamma, \varepsilon))(\mu^C - \mu^D l) + \alpha(\gamma, \varepsilon)V_{k+1}$, and

$$V_{k+1} = (1 - \gamma)\mu^C + \gamma(1 - \varepsilon)\mu^C V_{k+1} + \gamma(1 - (1 - \varepsilon)\mu^C)V_{k+2},$$

so $V_{k+1} = (1 - \beta(\gamma, \varepsilon, \mu^C))\mu^C + \beta(\gamma, \varepsilon, \mu^C)V_{k+2}$. Thus, we obtain

$$V_k = \alpha(\gamma, \varepsilon)(1 - \beta(\gamma, \varepsilon, \mu^C))\mu^C + (1 - \alpha(\gamma, \varepsilon))(\mu^C - \mu^D l) + \alpha(\gamma, \varepsilon)\beta(\gamma, \varepsilon, \mu^C)V_{k+2}.$$

If the player changed their strategy by playing according to P at record k and according to S at record $k + 1$, but otherwise kept the strategy the same, their expected continuation payoff at record $k + 1$, \tilde{V}_{k+1} , would be $\tilde{V}_{k+1} = (1 - \alpha(\gamma, \varepsilon))(\mu^C - \mu^D l) + \alpha(\gamma, \varepsilon)V_{k+2}$, and their expected payoff upon reaching record k , which we denote by \tilde{V}_k , would be

$$\begin{aligned} \tilde{V}_k &= (1 - \beta(\gamma, \varepsilon, \mu^C))\mu^C + \beta(\gamma, \varepsilon, \mu^C)\tilde{V}_{k+1} \\ &= (1 - \beta(\gamma, \varepsilon, \mu^C))\mu^C + (1 - \alpha(\gamma, \varepsilon))\beta(\gamma, \varepsilon, \mu^C)(\mu^C - \mu^D l) + \alpha(\gamma, \varepsilon)\beta(\gamma, \varepsilon, \mu^C)V_{k+2}. \end{aligned}$$

Then $\tilde{V}_k - V_k = (1 - \alpha(\gamma, \varepsilon))(1 - \beta(\gamma, \varepsilon, \mu^C))\mu^D l$. Note that $\mu^D > 0$, for if $\mu^D = 0$, then every record must be a cooperator record, in which case every player would always play D , which is a contradiction. Hence $\tilde{V}_k - V_k > 0$, and the profile is not an

equilibrium. ■

These lemmas show that only generalized trigger strategies can satisfy both forgery-proofness and coordination-proofness. For the converse, note that every generalized trigger strategy satisfies coordination-proofness. Moreover, if a generalized trigger strategy is an equilibrium, it must satisfy $V_k > V_{k+1}$ for all $k < K + L$, as otherwise a record k player would prefer to defect. Since V_k is constant for all $k \geq K + L$, V_k is everywhere non-increasing, so forgery-proofness is satisfied.