

Record-Keeping and Cooperation in Large Societies*

Daniel Clark, Drew Fudenberg, and Alexander Wolitzky

January 4, 2021

Abstract

We introduce a new model of repeated games in large populations with random matching, overlapping generations, and limited records of past play. We prove that steady-state equilibria exist under general conditions on records. When the updating of a player’s record can depend on the actions of both players in a match, any strictly individually rational action can be supported in a steady-state equilibrium. When record updates can depend only on a player’s own actions, fewer actions can be supported. Here we focus on the prisoner’s dilemma and restrict attention to strict equilibria that are *coordination-proof*, meaning that matched partners never play a Pareto-dominated Nash equilibrium in the one-shot game induced by their records and expected continuation payoffs. Such equilibria can support full cooperation if the stage game is either “strictly supermodular and mild” or “strongly supermodular,” and otherwise permit no cooperation at all. The presence of “supercooperator” records, where a player cooperates against any opponent, is crucial for supporting any cooperation when the stage game is “severe.”

Keywords: repeated games, community enforcement, steady-state equilibria, records, cooperation, prisoner’s dilemma

*This paper was previously distributed with the title “Steady-State Equilibria in Anonymous Repeated Games.” It replaces our earlier papers “Steady-State Equilibria in Anonymous Repeated Games, I: Trigger Strategies in General Stage Games,” “Steady-State Equilibria in Anonymous Repeated Games, II: Coordination-Proof Strategies in the Prisoner’s Dilemma,” and “Robust Cooperation with First-Order Information.” 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

In many settings of economic interest, individuals interact with different partners over time, and bad behavior against one partner causes a negative response by other members of society. Moreover, people often have fairly limited information about their partners' past behavior, and little to no information about the behavior of people with whom their partners previously interacted. Yet, groups often maintain outcomes that are more efficient than those consistent with myopic incentives.¹

To study these situations, we introduce a new class of repeated games with random matching with three key features. First, there is a continuum population, where individuals have geometrically distributed lifespans (with deaths balanced by a constant inflow of new players). Second, all that players know about each partner's past behavior or social standing is the partner's current *record*. Third, the time horizon is doubly infinite (so there is no commonly known start date or notion of calendar time), and we analyze *steady states* where the population distribution of records is constant over time. Compared to standard repeated game models with a fixed finite set of players, a commonly known start date, and a common notion of calendar time and/or a public randomizing device, our model seems more appropriate for studying cooperation in large decentralized societies. In addition, the combination of the continuum population and steady state assumptions keeps the model tractable even in the presence of recording or implementation errors, because individual agents do not learn about the state of the system from their own observations. On the other hand, a new challenge in our setting is managing the interaction of incentive conditions (which depend on the steady-state shares of players with different records) and steady-state conditions (which depend on equilibrium strategies).

Two fundamental questions about such an environment are “What sort of records is

¹Examples of such “community enforcement” or “indirect reciprocity” include Milgrom, North, and Weingast (1990) and Greif (1993) on merchant coalitions; Klein and Leffler (1981), Resnick and Zeckhauser (2002), and Dellarocas (2005) on seller reputation; Klein (1992) and Padilla and Pagano (2000) on credit ratings; and Friedman and Resnick (2001) on online ratings.

a society likely to generate?” and “What sorts of records suffice to support cooperative behavior?” Like most previous studies of record-keeping in community enforcement, this paper focuses exclusively on the second of these questions, briefly discussing the first in the conclusion. Our main finding is that in many settings supporting cooperation requires that records contain not only information about individuals’ past actions, but also information about the context of these actions. However, such contextualizing information need not be extremely detailed—it is enough to record how players’ immediate partners behaved towards them.

The record-keeping systems we study can be viewed as idealizations of the types of information that large societies need to support cooperation, but there are also some real-world settings where they can be taken more literally. Some examples are the online rating systems used by platforms like eBay, AirBnB, and Uber. There is strong evidence that users’ ratings on these platforms determine their prospects for finding future trading partners, even after controlling for unobserved heterogeneity (Cabral and Hortasçu (2010), Luca (2016)). On some platforms (e.g., eBay, Amazon Marketplace) users have no information about their current partner’s past partners’ behavior—these are examples of what we will call *first-order* systems. On other platforms (e.g., AirBnB), users can look up feedback that their partners left about their own past partners, a form of what we will call *second-order* information. While many considerations influence a platform’s choice of rating system (Tadelis, 2016), our model highlights the ability to distinguish justified and unjustified deviations from desired equilibrium behavior as a factor that favors systems capable of recording second-order information.

In our model, a *record system* updates the players’ records based on their current records and the actions they choose. These systems may be stochastic, due to either recording errors or errors in implementing a player’s intended action. We prove that steady states exist for record systems that are *finite partitional*, which means that for any record, there is a finite partition of the opponent’s record space such that the update function does not vary with the opponent’s record within a given partition

element. This condition is quite general; it is satisfied by all record systems we analyze as well as those considered in prior work.

We then characterize the prospects for steady-state cooperation under different types of record systems. To capture a simple form of robustness, we consider only strict equilibria. For most of our results, we also require equilibria to be *coordination-proof*, which means that a pair of matched players never play a Pareto-dominated equilibrium in the “augmented” game induced by their current records and their expected continuation payoffs. Restricting attention to coordination-proof strategies rules out equilibria built on within-match miscoordination.² Finally, our positive results focus on the double limit where players’ expected lifespans are long and there is little noise in the record system.³ Taking this limit allows a clean analysis and gives society its best shot at supporting cooperative outcomes; of course, if players are myopic or records are extremely noisy, only static Nash equilibrium outcomes can arise.

We begin by analyzing *second-order records*, where record updates depend on a player’s own record and action as well as their partner’s action, but not their partner’s record. In other words, a player’s second-order record depends only on their own past actions and their past partners’ actions towards them. We show that second-order records are rich enough to support the play of any action that Pareto-dominates the pure-strategy minmax payoff (in the long-lifespans, low-noise double limit).⁴ To prove this, we consider strategies that assign players to good or bad standing based on their records, and specify that good-standing players take the target action when matched with each other, while players take a minmaxing action whenever at least one player in the match has bad standing. With these strategies, second-order records can identify the standing of a good-standing player’s partner from their action. This allows the threat of switching a good-standing player to bad standing to incentivize

²Coordination-proofness is 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 a single one-shot game. This simplicity is another advantage of our model.

³These results also assume record systems are “canonical,” as defined in Section 2.5.

⁴The equilibria we construct to prove this “minmax-threat” folk theorem need not be coordination-proof. We also provide a “Nash-threat” folk theorem in coordination-proof equilibria.

the prescribed behavior among good-standing players; similarly, bad-standing players can be incentivized by the promise of an eventual return to good standing.

We then turn to *first-order records*, where a player’s record is updated based only on their own record and action (as in the eBay and Amazon Marketplace examples). First-order record systems cannot support as many actions as second-order systems, because first-order records lack the contextualizing information required to distinguish justified deviations from the target equilibrium action from unjustified ones. For example, in the prisoner’s dilemma (“PD”), if a player is penalized for defecting against opponents who cooperate, they must be equally penalized for defecting against opponents who defect. This impossibility of conditioning rewards and punishments on the opponent’s action can preclude steady-state cooperation.

We first highlight a type of situation where the inability to distinguish justified from unjustified deviations does *not* pose a major obstacle to supporting a target equilibrium action a . This occurs when there exists an “unprofitable punishment” action b such that a player is made worse-off when their partner switches from a to b , but unilaterally deviating to b is not profitable when the opponent is expected to play a .⁵ For example, in the PD, *Defect* is *not* an unprofitable punishment for *Cooperate* because it violates the second condition: unilaterally deviating to *Defect* is profitable when the opponent plays *Cooperate*. In settings where an unprofitable punishment for action a does exist, strategies based on first-order records can support the play of a by penalizing a player only if they take an action other than a or b . Intuitively, the inability to distinguish justified and unjustified plays of b is not an obstacle to supporting a , since no one has an incentive to unilaterally deviate to b .

Our positive results for second-order records and for first-order records with unprofitable punishments raise the question of when an action without an unprofitable punishment can be supported with first-order records. The remainder of our analy-

⁵There is also an additional, more subtle requirement: there must exist a best response c to b such that b is a better response to c than a is. We explain the role of this additional requirement in Section 4.

sis answers this question for the leading example of cooperation in the PD. That is, we characterize the set of payoff parameters in the PD for which there exist strict, coordination-proof equilibria in which the share of cooperation converges to 1 in the long-lifespans, low-noise double limit. The characterization is fairly subtle: we find that full limit cooperation is possible if either (i) the degree of strategic complementarity in the PD stage game is sufficiently high, or (ii) the degree of complementarity is positive and in addition the instantaneous gain from defection is sufficiently low; and that otherwise the only strict, coordination-proof equilibrium is *Always Defect*. Interestingly, cooperation in case (i) requires that a non-zero share of players have records at which the equilibrium strategy prescribes cooperation if and only if their opponent is going to cooperate, and also that a non-zero share of players have records at which cooperation is prescribed even if the opponent is going to defect. The latter *supercooperator* records prevent too many players from transiting to “bad” records where they always defect.

There is a small prior literature on record-keeping in community enforcement. Without noise, Okuno-Fujiwara and Postlewaite (1995) established a folk theorem with interdependent records (which are more permissive than our second-order records).⁶ Takahashi (2010) constructed efficient equilibria in the PD when players observe their partner’s entire history of actions—all first-order information—but no higher-order information. That paper did not consider steady states, so it did not contend with the interaction of incentive and steady-state conditions, and its conditions for efficient equilibria to exist are more permissive than ours. Heller and Mohlin (2018) constructed an efficient mixed-strategy equilibrium in a PD where players are completely patient and observe a finite sample of their current partner’s past actions, and a small fraction of

⁶Sugden (1986) and Kandori (1992) proved related results. Antecedents include Rosenthal (1979) and Rosenthal and Landau (1979), which focused on existence results and examples. Steady-state equilibria in models with interdependent records also appear in the literature on fiat money (e.g., Kiyotaki and Wright (1993), Kocherlakota (1998)). A less closely related literature studies community enforcement in finite populations without any information beyond the outcome of one’s own matches (e.g., Kandori (1992), Ellison (1994), Dilmé (2016), Deb, Sugaya, and Wolitzky (2020)). With so little information, cooperation cannot be supported in a continuum population, or in a finite population that is large compared to the discount factor.

players are commitment types who always *Defect*.⁷

2 Framework

We consider a discrete-time random matching model with a constant unit mass of players, each of whom has a geometrically-distributed lifetime with continuation probability $\gamma \in (0, 1)$ (with exits balanced by a steady inflow of new entrants of size $1 - \gamma$). The time horizon is doubly infinite. When two players match, they play a finite, symmetric stage game with action space A and payoff function $u : A \times A \rightarrow \mathbb{R}$.

2.1 Record Systems

Every player carries a *record*, and when two players meet, each observes the other's record but no further information. Each player's record is updated at the end of every period in a “decentralized” way that depends only on their own action and record and their current partner's action and record.

Definition 1. A *record system* \mathcal{R} is a pair (R, ρ) comprised of a countable set R (the **record set**) and a function $\rho : R^2 \times A^2 \rightarrow \Delta(R)$ (the **update rule**), where the four arguments of ρ are (in order) a player's own record, current partner's record, own action, and current partner's action.

Note that the update rule is allowed to be stochastic. This can capture errors in recording, as well as imperfect implementation of players' intended actions.⁸ We assume all newborn players have the same record, which we denote by 0. (Our main results extend to the case of a non-degenerate, exogenous distribution over initial records.)

⁷Heller and Mohlin (2018) also considered information structures that are similar to our second-order records, but do not yield a folk theorem. Bhaskar and Thomas (2019) studied first-order information in a sequential-move game.

⁸In the imperfect implementation interpretation, the stage game payoffs are the expected payoffs that result when players intend to take the given stage game actions, and we hold these expected payoffs fixed as the noise level varies.

An update rule thus specifies a probability distribution over records as a function of a player’s record and action and their current partner’s record and action. We sometimes refer to the general case where ρ is unrestricted as an *interdependent* record system. An interdependent record system is *finite-partitional* if for each $r \in R$ there exists a finite partition $\bigcup_{m=1, \dots, M(r)} R_m = R$ such that, whenever $r', r'' \in R_m$ for some m , $\rho(r, r', a, a') = \rho(r, r'', a, a')$ for all $a, a' \in A$. Kandori (1992)’s “local information processing” and Okuno-Fujiwara and Postlewaite (1995)’s “status levels” are examples of finite-partitional interdependent record systems.

Many simple and realistic record systems fall into a more restricted class, where a player’s record update does not depend directly on their opponent’s record.

Definition 2. *A record system is **second-order** if the update rule can depend only on a player’s own action and record and their partner’s action, i.e. $\rho(r, r', a, a') = \rho(r, r'', a, a')$ for all $r, r', r'' \in R, a, a' \in A$.*

In the absence of noise, with a second-order record system, a player’s record can be computed based only on their own history of stage-game outcomes.

Finally, in some situations a player’s record depends only on their own actions.

Definition 3. *A record system is **first-order** if the update rule can depend only on a player’s own action and record, i.e. $\rho(r, r', a, a') = \rho(r, r'', a, a'')$ for all $r, r', r'' \in R, a, a', a'' \in A$.*

Nowak and Sigmund (1998), Panchanathan and Boyd (2003), Takahashi (2010), Bhaskar and Thomas (2019), and Heller and Mohlin (2018) also considered first-order records.⁹ We consider second-order records in Section 3 and first-order records in Sections 4 and 5. Note that both of these types of record system are finite-partitional.

⁹To interpret noisy first-order records as resulting from implementation errors, the outcome of the game must have a product structure in the sense of Fudenberg, Levine, and Maskin (1994), so that a player’s record update does not depend on the opponent’s action.

2.2 Strategies, States, and Steady States

In principle, each player can condition their play on the entire sequence of outcomes and past opponent records that they have observed. However, this information is payoff-irrelevant in a continuum-population steady state: the only payoff-relevant information available to a player is their own record and their current partner's record.

Thus, all strategies that condition only on payoff-relevant variables are *pairwise-public*, meaning that they condition only on information that is public knowledge between the two partners, namely their records. We restrict attention to such strategies. We write a pairwise-public pure strategy as a function $s : R \times R \rightarrow A$, with the convention that the first coordinate is the player's own record and the second coordinate is the partner's record, and similarly write a pairwise-public mixed strategy as a function $\sigma : R \times R \rightarrow \Delta(A)$. We also assume that all players use the same strategy.¹⁰

The *state* $\mu \in \Delta(R)$ of the system is the share of players with each possible record. To operationalize random matching in a continuum population, we specify that, when the current state is μ , the distribution of matches is given by $\mu \times \mu$. That is, 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 \mathcal{R} and a pairwise-public strategy σ , denote the distribution over next-period records of a player with record r who meets a player with record r' by $\phi_{r,r'}(\sigma) = \sum_a \sum_{a'} \sigma(r, r')[a]\sigma(r', r)[a']\rho(r, r', a, a') \in \Delta(R)$. Because newborn players have record 0, the *state update map* $f_\sigma : \Delta(R) \rightarrow \Delta(R)$ is given by

$$f_\sigma(\mu)[0] := 1 - \gamma + \gamma \sum_{r'} \sum_{r''} \mu_{r'} \mu_{r''} \phi_{r',r''}(\sigma)[0],$$

$$f_\sigma(\mu)[r] := \gamma \sum_{r'} \sum_{r''} \mu_{r'} \mu_{r''} \phi_{r',r''}(\sigma)[r] \text{ for } r \neq 0.$$

A *steady state* under strategy σ is a state μ such that $f_\sigma(\mu) = \mu$.

¹⁰As will become clear, restricting players to use the same, pairwise-public strategy is without loss for strict equilibria.

Theorem 1.

- (i) *Under any finite-partitional record system (and hence any first-order or second-order record system) and any pairwise-public strategy, a steady state exists.*
- (ii) *For record systems that are not finite-partitional, a steady state may fail to exist.*

The proof is in Appendix A.1. (All other omitted proofs can also be found in either the Appendix (A) or the Online Appendix (OA).) Intuitively, the combination of the finite domain of the record-update function (due to finiteness of the stage game and the finite-partition property) and geometrically distributed lifetimes imply that almost all players' records lie in a finite subset of the record set. This lets us find a convex set \bar{M} that contains all feasible distributions over records and resembles a finite-dimensional set—in particular, \bar{M} is compact in the sup norm. We then show that f maps \bar{M} to itself and is continuous in the sup norm so, since \bar{M} is also convex, we can appeal to a fixed point theorem.¹¹ When instead the record-update function does not have a finite domain, the update map can shift weight to the upper tail of the record distribution in such a way that no steady state exists. The proof shows that this is the case if $\gamma > 1/2$ and, whenever players with records r and r' meet, both of their records update to $\max\{r, r'\} + 1$.

Note that Theorem 1 does not assert that the steady state for a given strategy is unique, and it is easy to construct examples where it is not.¹² Intuitively, this multiplicity corresponds to different initial conditions at time $t = -\infty$.

¹¹Fudenberg and He (2018) used a similar proof technique. In that paper players do not observe each other's records, so the finite-partition property is automatically satisfied.

¹²For instance, suppose that $R = \{0, 1, 2\}$, the action set is singleton, and newborn players have record 0. When matched with a player with record 0 or 1, the record of a player with record 0 or 1 increases by 1 with probability ε and remains constant with probability $1 - \varepsilon$, but it increases by 1 with probability 1 when the player is matched with a player with record 2. When a player's record reaches 2, it remains 2 for the remainder of their lifetime. Depending on the parameters γ and ε , there can be between one and three steady states in this example.

2.3 Steady-State Equilibria

We focus on steady states that derive from equilibrium play. Given a record system \mathcal{R} , strategy σ , and state μ , define the flow payoff of a player with record r as

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

Denote the probability that a player with record r today has record r' in t periods if they are still alive by $\phi_r^t(\sigma, \mu)[r']$.¹³ The continuation value of a player with record r is

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

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

Each player's objective is to maximize their expected undiscounted lifetime payoff. Thus, a pair (σ, μ) is an *equilibrium* if μ is a steady state under σ and, for each own record r , opponent's record r' , and action a such that $\sigma(r, r')[a] > 0$, we have

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

An equilibrium is *strict* if the argmax is unique for all pairs of records (r, r') , so each player has a strict preference for following the equilibrium strategy. As noted earlier, every strict equilibrium is pairwise-public, pure, and symmetric. To distinguish equilibria (σ, μ) from Nash equilibria of the stage game, we call the latter *static equilibria*.

Corollary 1. *Under any finite-partitional record system, an equilibrium exists.*

Proof. Fix a symmetric static equilibrium α^* , and let σ recommend α^* at every record pair (r, r') . Then (σ, μ) is an equilibrium for any steady state μ . ■

¹³This is defined recursively by $\phi_r^1(\sigma, \mu)[r'] = \sum_{r''} \mu_{r''} \phi_{r, r''}(\sigma)[r']$ and for $t > 1$, $\phi_r^t(\sigma, \mu)[r'] = \sum_{r''} (\phi_r^{t-1}(\sigma, \mu)[r'']) (\phi_{r'', r}^1(\sigma, \mu)[r'])$.

Strict equilibria need not exist without additional assumptions; one sufficient condition is that the stage game has a strict and symmetric Nash equilibrium.

Corollary 2. *Under any finite-partitional record system, a strict equilibrium exists if the stage game has a strict and symmetric Nash equilibrium.*

The proof of Corollary 2 is identical to that of Corollary 1, except α^* is taken to be a strict and symmetric static equilibrium.

2.4 Coordination-Proofness

Coordination-proofness captures the idea that equilibria that rely on “miscoordination” within a match will break down if matched partners manage to coordinate successfully. For a fixed equilibrium (σ, μ) , denote the expected continuation payoff of a player with record r who plays action a against an opponent with record r' who plays a' by $V_{r,r'}^{a,a'} := \sum_{r''} \rho(r, r', a, a', r'') V_{r''}$. The *augmented payoff function* $\hat{u} : R \times R \times A \times A \rightarrow \mathbb{R}$ is defined by $\hat{u}_{r,r'}(a, a') := (1 - \gamma)u(a, a') + \gamma V_{r,r'}^{a,a'}$. The *augmented game* between players with records r and r' is the static game with action set $A \times A$ and payoff functions $\hat{u}_{r,r'}$ and $\hat{u}_{r',r}$.

Since (σ, μ) is an equilibrium, the prescribed stage-game strategy profile $(\sigma(r, r'), \sigma(r', r))$ is a Nash equilibrium in the augmented game between players with records r and r' for any $(r, r') \in R^2$. We say that the equilibrium is *coordination-proof* if $(\sigma(r, r'), \sigma(r', r))$ is never Pareto-dominated by another augmented-game Nash equilibrium.

Definition 4. *An equilibrium (σ, μ) is **coordination-proof** if, for any records r, r' and any Nash equilibrium (α, α') in the augmented game between players with records r and r' , if $\hat{u}_{r,r'}(\alpha, \alpha') > \hat{u}_{r,r'}(\sigma(r, r'), \sigma(r', r))$ then $\hat{u}_{r',r}(\alpha', \alpha) < \hat{u}_{r',r}(\sigma(r', r), \sigma(r, r'))$.*

The logic is that, if $(\sigma(r, r'), \sigma(r', r))$ were Pareto-dominated by some augmented-game Nash equilibrium (α, α') , players with records r and r' would benefit from reaching a self-enforcing agreement to play (α, α') when matched together, breaking the equilibrium. The following corollary (proved in OA.1) gives a sufficient condition for the

existence of a coordination-proof equilibrium that is satisfied in many games, including the PD.

Corollary 3. *Under any finite-partitional record system, a coordination-proof equilibrium exists if the stage game has a symmetric Nash equilibrium that is not Pareto-dominated by another (possibly asymmetric) Nash equilibrium.*

2.5 Canonical Records

For our positive results, we focus on canonical record systems. These track either the sequence of actions (for first-order records) or stage-game outcomes (for second-order records) in a player's history, and allow for possible misrecording in any period.

Let $n = |A|$. With second-order records, a *noise matrix* ε is an $n^2 \times n^2$ matrix with diagonal elements equal to 0 and non-negative off-diagonal elements, where $\varepsilon_{(a,a'),(\tilde{a},\tilde{a}'})$ is the probability that outcome (a, a') is mis-recorded as $(\tilde{a}, \tilde{a}') \neq (a, a')$. The canonical second-order record set is the set of finite sequences of pairs of actions, $\bigcup_{t=0}^{\infty} (A \times A)^t$. Given a second-order canonical record $r = \prod_{\tau=1}^t (a_{\tau}, a'_{\tau})$ and an outcome (a, a') , $(r, (a, a'))$ is the canonical record formed by concatenating r and (a, a') .

Definition 5. *A second-order record system is **canonical** if the record set R is canonical and there exists a noise matrix ε such that, for every record $r = \prod_{\tau=1}^t (a_{\tau}, a'_{\tau})$ and action pair (a, a') , we have*

$$\rho(r, a, a') = \left(1 - \sum_{(\tilde{a}, \tilde{a}') \neq (a, a')} \varepsilon_{(a, a'), (\tilde{a}, \tilde{a}')} \right) \rho(r, (a, a')) + \sum_{(\tilde{a}, \tilde{a}') \neq (a, a')} \varepsilon_{(a, a'), (\tilde{a}, \tilde{a}')} \rho(r, (\tilde{a}, \tilde{a}')).$$

Similarly, in a canonical first-order record systems, records are sequences of actions and each action a has probability $\varepsilon_{a, \tilde{a}}$ of being misrecorded as action $\tilde{a} \neq a$.

In general, the set of equilibria depends on both the amount of noise in the system and the players' expected lifetimes. We focus on the case where there is little or no noise, and players live a long time. We thus consider the double limit $(\gamma, \varepsilon) \rightarrow (1, 0)$,

where ε is the noise matrix in a canonical record system, and $\varepsilon \rightarrow 0$ means that every entry in the matrix ε converges to 0.

Definition 6. Let $\bar{\mu}^a(\gamma, \varepsilon)$ denote the supremum of the share of players taking action a over all equilibria for parameters (γ, ε) . Action a is **limit-supported** if $\lim_{(\gamma, \varepsilon) \rightarrow (1, 0)} \bar{\mu}^a(\gamma, \varepsilon) = 1$.

3 Second-Order Records: Folk Theorems

Our next result shows that a wide range of actions can be limit-supported with second-order records. Because second-order records allow a player's record update to depend on both players' actions, we can construct strategies that punish opportunistic actions but avoid punishing players who punish others when they are supposed to. For example, in the prisoner's dilemma our strategies count *Defect vs. Cooperate* as a "bad" outcome, but not *Defect vs. Defect*, a distinction that cannot be made using first-order records.

Denote the pure-strategy minmax payoff by $\underline{u} = \min_{a'} \max_a u(a, a')$.

Theorem 2. Fix an action a . With canonical second-order records:

- (i) If $u(a, a) > \underline{u}$, then a can be limit-supported by strict equilibria.
- (ii) If $u(a, a) > u(b, b)$ for some strict and symmetric static equilibrium (b, b) that is not Pareto-dominated by another static equilibrium, then a can be limit-supported by strict, coordination-proof equilibria.

Theorem 2(i) is a minmax-threat folk theorem. The construction relies on "cyclic" strategies of the following form: Let $b \in \arg \min_{a'} \max_a u(a, a')$ be a minmax action. Players begin in good standing. A player in good standing plays a when matched with fellow good-standing players and plays b against bad-standing players, while a player in bad standing always plays b . When a good-standing player's outcome is recorded as anything other than (a, a) or (b, b) , the player enters bad standing. A player remains in bad standing until they accumulate M (b, b) profiles for some fixed $M \in \mathbb{N}$, at

which point they return to good standing. We show that, when $(\gamma, \varepsilon) \approx (1, 0)$, M can be chosen to be high enough so that the punishment of M periods of (b, b) is severe enough to deter deviations from the prescribed strategy, but also low enough that the steady-state share of players in good standing is high.

This equilibrium may not be coordination-proof. For example, suppose there is a symmetric static equilibrium (c, c) such that $u(c, c)$ is significantly greater than $u(a, a)$. Then a pair of bad-standing players may benefit from reaching a self-enforcing agreement to play (c, c) rather than (b, b) , even though this delays their return to good standing by one period.

Theorem 2(ii) presents a condition under which an action a can be limit-supported by strict, coordination-proof equilibria. It gives a Nash-threat folk theorem, where the “threat point” equilibrium (b, b) is required to be strict, symmetric, and not Pareto-dominated by another static equilibrium. For example, in the prisoner’s dilemma, taking $a = C$ and $b = D$ implies that *Cooperate* is limit-supported by strict, coordination-proof equilibria.

The proof of part (ii) uses the following grim trigger strategy. A player whose outcome has never been recorded as anything other than (a, a) or (b, b) is in good standing, and all other players are in bad standing. Players in good standing play a against fellow good-standing players and play b against bad-standing players, while bad-standing players always play b . Such strategies can support cooperation in the iterated limit where first noise becomes small ($\varepsilon \rightarrow 0$) and then players become long-lived ($\gamma \rightarrow 1$). To handle the general limit, we modify this strategy with an adaptation of the threading technique used in papers such as Ellison (1994) and Takahashi (2010).¹⁴ In particular, for a given $N \in \mathbb{N}$, a pair of matched players condition their play only on the the recordings of each other’s outcomes in periods which preceded the current period by a multiple of N . Thus, within each thread, the effective continuation probability

¹⁴We will use this threading technique in many of our results. In unpublished earlier work (Clark, Fudenberg, and Wolitzky, 2019a,b) we consider simpler strategies that deliver similar but less clean results without using threads.

is γ^N rather than γ . By altering N , we are able to control the effective continuation probability, and can essentially replicate the iterated limit where first noise becomes small and then players become long-lived.

Okuno-Fujiwara and Postlewaite (1995)’s Theorem 1 showed that with a form of interdependent records (termed “status”), any actions that Pareto-dominate the pure-strategy minmax payoffs can be supported without noise. Their proof uses grim trigger strategies, so it not robust to noise. Theorem 2 shows that their theorem’s conclusion does not require interdependent records and also extends to settings with overlapping generations and noise.

4 First-Order Records: Unprofitable Punishments

Now we turn to first-order record systems, where a player’s record depends only on their own past play. Such records cannot support as many actions as second-order records can, and the folk theorem fails for strict equilibrium. The key obstacle is that first-order records cannot distinguish “justified” deviations from the target action profile from “unjustified” ones. For example, in the PD, if players are penalized for playing *Defect* against *Cooperate* (an off-path, opportunistic deviation), they must be equally penalized for playing *Defect* against *Defect* (a justified punishment that must sometimes occur on-path if defection is to be deterred). As we will see, this obstacle precludes cooperation in some games.

This section shows that this obstacle does not arise when the target action profile (a, a) has the property that there exists a punishing action b and a strict best response c to b such that $u(a, a) > u(c, b)$ (so that facing b is indeed a punishment), $u(a, a) > u(b, a)$ (so that deviating from a to b is unprofitable for a player whose opponent plays a), and $u(b, c) > u(a, c)$ (so a player prefers to carry out the punishment b rather than play the target action when the opponent best-responds to b). We say that in this case b is an *unprofitable punishment* for a . Intuitively, when an unprofitable punishment b exists for action a , the threat of switching to b can motivate one’s opponents to play a ,

but a player is not tempted to unilaterally deviate to b against opponents who play a . This enables first-order records to support the play of (a, a) by penalizing players only for taking actions other than a or b . In contrast, when the only punishing action b is a tempting deviation against a (as in the PD, where the punishing action D is always tempting), players must be penalized for playing b , and the record system's inability to distinguish justified and unjustified plays of b becomes a real obstacle.

Theorem 3. *Fix an action a . With canonical first-order records, if there exists an unprofitable punishment b for a and there is a strict and symmetric static equilibrium (d, d) , then a can be limit-supported by strict equilibria.*

The proof of Theorem 3, which is in OA.2, is similar to the proof of Theorem 2(ii), except now a player transitions to bad standing whenever their action is recorded as anything other than a or b (rather than transitioning whenever their action *profile* is recorded as anything other than (a, a) or (b, b)), and bad-standing players play c against good-standing players but play d against fellow bad-standing players.¹⁵

Setting $b = c = d$ in Theorem 3 yields the following corollary.

Corollary 4. *If there exists an action b such that (b, b) is a strict static equilibrium and $u(a, a) > \max\{u(b, a), u(b, b)\}$, then a can be limit-supported by strict equilibria.*

Here the strategy used to prove Theorem 3 takes the form of grim trigger with reversion to (b, b) . For example, suppose the stage game is a PD with an *exit option* E , which gives both players an “exit payoff” that is less than the cooperative payoff $u(C, C)$ whenever either of them plays E . Then (E, E) is a static equilibrium, and E is not a profitable deviation against C , so cooperation can be limit-supported by strategies that revert to (E, E) .¹⁶

¹⁵Another difference is that the equilibria used to prove Theorem 3 may not be coordination-proof. This is because there may be some static best response to a , e , such that (a, e) is a Nash equilibrium in the augmented game between a good-standing and bad-standing player that Pareto-dominates the prescribed action profile (b, c) .

¹⁶Technically, to apply Theorem 3 we need to specify that $u(E, E) > \max\{u(C, E), U(D, E)\}$. If instead this inequality is replaced with an equality, the same proof applies, but the constructed equilibria are no longer strict.

This example relates to a debate regarding the role of punishment in the evolution of human cooperation. The difficulty in distinguishing a warranted punishment from an unwarranted deviation is one factor that led Boyd et al. (2003) and Bowles and Gintis (2011) (among others) to argue that the enforcement of human cooperation cannot be explained without appealing to social preferences. Others (e.g., Baumard (2010), Guala (2012)) argued that human cooperation is better explained by simply avoiding deviators, rather than actively punishing them. The fact that cooperation in the PD is always limit-supported with second-order records, but (as we will see) is limit-supported with first-order records only for certain parameters, supports the argument that the inability to distinguish justified and unjustified plays of *Defect* is a serious obstacle to cooperation in the PD. However, this obstacle evaporates when a simple exit option is added to the game, consistent with Baumard and Guala’s position.

Another important example of unprofitable punishment arises when players can observably reduce their own utility by any amount while taking a stage-game action. In this case, whenever $0 < u(b, a) - u(a, a) < u(b, c) - u(a, c)$, the action “play b and burn some amount of utility in between $u(b, a) - u(a, a)$ and $u(b, c) - u(a, c)$ ” is an unprofitable punishment. That is, whenever the gain from playing b rather than a is greater when the opponent plays c as opposed to a , there exists an appropriate amount of utility that can be sacrificed to make playing b unattractive.

5 First-Order Records: Cooperation in the PD

In general stage games, characterizing whether an action that lacks an unprofitable punishment can be limit-supported with first-order records is challenging. In this section, we resolve this question for the leading case of cooperation in the PD. That is, we characterize the set of payoff parameters for which *Cooperate* can be limit-supported by strict, coordination-proof strategies. We will use the standard normalization of the PD payoffs, where the payoff to mutual cooperation is 1, the payoff to joint defection is 0, and the gain g from defecting against an opponent who cooperates and the loss

l from cooperating against an opponent who defects satisfy $g, l > 0$ and $l > g - 1$ (so that mutual cooperation 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

We first introduce some preliminary concepts in Sections 5.1 and 5.2. We then present our main characterization result in Section 5.3.

5.1 Defectors, Supercooperators, Preciprocators

We begin with some terminology for different types of records.

Definition 7. *Given a pure-strategy equilibrium (s, μ) , record r is a*

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

Defectors play D against all partners, while supercooperators play C against all partners, even those who will play D against them. In contrast, preciprocators exhibit a form of anticipatory reciprocation: they play C with partners whom they expect to play C , but play D with partners whom they expect to play D .

The PD is *strictly supermodular* if $g < l$, so the benefit of defecting is greater when the opponent defects, and conversely it is *strictly submodular* when $g > l$. A leading example of the PD 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 their opponent's play, 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. *Fix any first-order record system. 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, a supercooperator, or a reciprocator.*

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

1. Suppose that $g \geq l$. When two players with the same record r meet, by symmetry (an implication of strictness) 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 hence is also the strict best response to C , so record r is a defector.

2. When $g < l$, if 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 4. *Fix any first-order record system. If $g \geq l$, the unique strict equilibrium is Always Defect: $s(r, r') = D$ for all $r, r' \in R$.¹⁷*

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. ■

¹⁷The conclusion of Theorem 4 extends to all (possibly non-strict) pure-strategy equilibria whenever $g > l$. Takahashi (2010) and Heller and Mohlin (2018) obtained the same conclusion (as well as converse results) in related models.

An intuition for Theorem 4 is that a player’s continuation payoff decreases by the same amount whenever they play D , so they are willing to play D against opponents who play D while playing C against opponents who play C iff $g < l$. The theorem confirms that strictly individually rational actions are not always limit-supportable by strict equilibria with first-order records, in contrast to the situation with second-order records. Since this theorem resolves the submodular case, the rest of this section assumes strict supermodularity: $g < l$.

5.2 Coordination-Proofness in the Supermodular PD

We note some simple consequences of coordination-proofness in the strictly supermodular PD.

Lemma 2. *Fix any first-order record system. In any strict, coordination-proof equilibrium in the strictly supermodular PD, whenever two preciprocators meet, they play (C, C) .*

Proof. By definition, preciprocators play C against opponents who play C and play D against those who play D . Hence, the augmented game between any two preciprocators is a coordination game, with Nash equilibria (C, C) and (D, D) . 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. Therefore, the (C, C) equilibrium yields both higher stage-game payoffs and higher continuation payoffs than the (D, D) equilibrium, and thus Pareto-dominates it. So coordination-proofness dictates that any pair of matched preciprocators must play (C, C) rather than (D, D) . ■

Coordination-proofness thus implies that every preciprocator plays C when matched with another preciprocator or a supercooperator, and plays D when matched with a defector. In particular, all preciprocators play C against the same set of opposing records. Hence, the strategy in a strict, coordination-proof equilibrium is completely characterized by a description of which records are preciprocators, which are supercooperators, and which are defectors. Denote the total population shares of these records

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 $\mu^C = \mu^P + \mu^S = 1 - \mu^D$.

5.3 Limit-Supporting Cooperation in Supermodular PD's

We now present necessary and sufficient conditions for cooperation to be limit-supported in strict, coordination-proof equilibria with first-order records. Our sufficient conditions require canonical records with $\varepsilon \rightarrow 0$, while our necessary conditions apply for any “noisy” first-order record system.

Definition 8. *A first-order record system is **noisy** if for each record r there exist $q_C(r), q_D(r) \in \Delta(R)$ and $\varepsilon_C(r) \in (0, 1/2]$, $\varepsilon_D(r) \in [0, 1/2]$ such that*

$$\begin{aligned}\rho(r, C) &= (1 - \varepsilon_C(r))q_C(r) + \varepsilon_C(r)q_D(r), \text{ and} \\ \rho(r, D) &= \varepsilon_D(r)q_C(r) + (1 - \varepsilon_D(r))q_D(r).\end{aligned}$$

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 ε 's represent noise. The key feature of this definition is that perfect recording of actions is ruled out by the assumption that $\varepsilon_C(r) > 0$.

We say that the prisoner's dilemma is *mild* if $g < 1$ and *severe* otherwise, and that the game is *strongly supermodular* if $l > g + g^2$.

Theorem 5. *In the strictly supermodular PD,*

- (i) *With any noisy first-order 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 strict, coordination-proof equilibrium is Always Defect.*
- (ii) *With canonical first-order records, if either $g < 1$ or $l > g + g^2$, cooperation can be limit-supported by strict, coordination-proof equilibria.*




Figure 2: Limit efficiency obtains in the blue region. In the red region, the only strict, coordination-proof equilibrium is *Always Defect*.

Figure 2 displays the conclusions of Theorem 5. 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 2.

We now discuss the intuition for the necessary and sufficient conditions in Theorem 5. The proofs are contained in A.3.

5.3.1 Necessary Conditions

Broadly speaking, a small value of g makes supporting cooperation easier by reducing a reciprocator's temptation to deviate to D against an opponent who is expected to play C , while a large value of l makes supporting cooperation easier by reducing a reciprocator's temptation to deviate to C against an opponent who is expected to play D . The specific necessary condition $g < 1$ or $l > g + g^2$ comes from combining two inequalities: $\mu^S < 1/(1 + g)$ and $\mu^P + \mu^S(l - g) > g$. Note that the latter inequality requires that $\mu^S > 0$ when $g \geq 1$: in a severe prisoner's dilemma, there must be a positive share of supercooperators in any strict, coordination-proof equilibrium with any cooperation at all. The next lemma shows that combining these inequalities

delivers the necessary condition $g < 1$ or $l > g + g^2$. After the lemma's short proof, we explain why the inequalities hold.

Lemma 3. *If $g \geq 1$ and $l \leq g + g^2$, it is not possible to satisfy both $\mu^S < 1/(1 + g)$ and $\mu^P + \mu^S(l - g) > g$.*

Proof. Suppose that $\mu^S < 1/(1 + g)$. Then $\mu^P + \mu^S(l - g)$ 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 $\mu^S = 1/(1 + g)$. Hence, if $\mu^P + \mu^S(l - g) > g \geq 1$, then $l/(1 + g) > g$, which requires $l > g + g^2$. \blacksquare

To derive the inequality $\mu^S < 1/(1 + g)$, note that a defector's flow payoff equals $\mu^S(1 + g)$, as defectors receive payoff $1 + g$ when matched with supercooperators, and otherwise receive payoff 0. This flow payoff must be less than 1, since otherwise it would be optimal for newborn players to play D for their entire lives instead of following the equilibrium strategy.

The inequality $\mu^P + \mu^S(l - g) > g$ is established by Lemma 9 in A.3.2.1. As this inequality is a key point where incentive conditions and steady-state conditions come together to determine the scope of cooperation, here we provide a derivation for the special case where there is a "best" record $r^* = \arg \max_r V_r$, r^* is a preciprocator record, and the noise levels $\varepsilon_C(r) = \varepsilon_C$ and $\varepsilon_D(r) = \varepsilon_D$ are independent of r . Since every preciprocator has an expected flow payoff of μ^C and the probability that a preciprocator is recorded as playing C is $(1 - \varepsilon_C)\mu^C + \varepsilon_D\mu^D$, we have

$$V_{r^*} = (1 - \gamma)\mu^C + \gamma((1 - \varepsilon_C)\mu^C + \varepsilon_D\mu^D)V_{r^*}^C + \gamma(\varepsilon_C\mu^C + (1 - \varepsilon_D)\mu^D)V_{r^*}^D,$$

where $V_{r^*}^C$ and $V_{r^*}^D$ denote the expected continuation payoffs of record r^* players who are recorded as playing C and D , respectively. The incentive constraint for a record r^* player to play C against an opponent who plays C is $(1 - \gamma)(1) + \gamma(1 - \varepsilon_C)V_{r^*}^C + \gamma\varepsilon_C V_{r^*}^D > (1 - \gamma)(1 + g) + \gamma\varepsilon_D V_{r^*}^C + \gamma(1 - \varepsilon_D)V_{r^*}^D$, or $\gamma(1 - \varepsilon_C - \varepsilon_D)(V_{r^*}^C - V_{r^*}^D)/(1 - \gamma) > g$. By

the accounting identity above, this simplifies to

$$\frac{1 - \varepsilon_C - \varepsilon_D}{\varepsilon_C \mu^C + (1 - \varepsilon_D) \mu^D} \left(\mu^C - V_{r^*} - \frac{\gamma}{1 - \gamma} (V_{r^*} - V_{r^*}^C) \right) > g.$$

Now note that in a steady state the expected lifetime payoff of a newborn player equals the average flow payoff in the population in a given period: $V_0 = \mu^P \mu^C + \mu^S (\mu^C - \mu^D l) + \mu^D \mu^S (1 + g)$. Since $V_{r^*} \geq V_0$ and $V_{r^*} \geq V_{r^*}^C$, the incentive constraint implies

$$\frac{(1 - \varepsilon_C - \varepsilon_D) \mu^D}{\varepsilon_C \mu^C + (1 - \varepsilon_D) \mu^D} (\mu^P + \mu^S (l - g)) > g,$$

which itself implies $\mu^P + \mu^S (l - g) > g$.

The necessary conditions for cooperation in Theorem 5(i) hold for any noisy first-order record system. The reason for assuming noise is that the proof analyzes incentives at cooperator records where there is a positive probability of being recorded as playing D in a given period. Without noise, there can be fully cooperative equilibria even when $g \geq 1$ and $l \leq g + g^2$. For example, for some parameters grim trigger strategies, together with the steady state where no one has ever defected, form an equilibrium.

5.3.2 Sufficient Conditions

We use different types of strategies to support cooperation when $g < 1$ and when $g \geq 1$ and $l > g + g^2$. The strategies used in the $g < 1$ case are threaded grim trigger strategies, similar to those used to prove Theorem 2(ii) and Theorem 3, as well as Proposition 1 in Takahashi (2010).

When $g \geq 1$, we have seen that cooperation requires supercooperator records: in particular, grim trigger strategies cannot support cooperation. Consider instead strategies that take the following form for some J, K : Within each thread, players begin life as defectors, become reciprocators once they have been recorded as playing D J times, transition to being supercooperators once they have been recorded as playing D an additional K times; and finally permanently transition to being defectors

once they have been recorded as playing D once more. Two features of the resulting equilibria are particularly notable.

First, players' incentives to cooperate are provided solely by the threat of entering defector status once the number of times they have been recorded as playing D (their "score") reaches $J + K + 1$, at which point reciprocator opponents switch from playing C to D against them. (In contrast, when a player switches from reciprocator to supercooperator status, their opponents' behavior is unaffected.) Since the survival probability γ is less than 1, this threat looms larger the closer a player's score is to $J + K + 1$. Hence, players with higher scores are willing to incur greater costs to prevent their scores from increasing further. Our construction exploits this observation by finding a critical score $J + K$ such that players with score $J + K$ are willing to play C at a cost of l , while players with scores less than K are willing to play C at a cost of g but not at a cost of l . That is, players with score $J + K$ supercooperate, while those with scores from J to $J + K - 1$ reciprocate.

Second, the feature that players with score $J + K$ supercooperate rather than reciprocate may at first seem to work against cooperation, because defectors obtain higher payoffs against supercooperators than cooperators do. However, the presence of supercooperators increases the steady-state share of reciprocators, via the following mechanism: Since players with score $J + K$ supercooperate, their scores increase more slowly than if they reciprocated. Therefore, fewer players survive to enter defector status, which reduces the steady-state share of defectors. Finally, when there are fewer defectors, reciprocators defect less often, and hence their scores increase more slowly, which increases the steady-state share of reciprocators. In sum, the presence of supercooperators reduces the steady-state share of defectors and increases the steady-state share of reciprocators, which enables steady-state cooperation.¹⁸

¹⁸Ostrom (1990) found that giving norm violators opportunities to resume cooperation before facing harsher punishments helps sustain cooperation by preventing excessively fast breakdowns following occasional violations. The mild punishment of transitioning to supercooperator status serves a broadly similar role in our model.

6 Discussion

This paper introduces a new model of repeated social interactions, where players interact with a sequence of anonymous and random opponents, and their information about their opponents' past play consists of noisy "records." We study steady-state equilibria in a large population with geometrically distributed lifetimes, focusing on situations where there is little noise and lifetimes are long.

We find that any strictly individually rational outcome can be supported with second-order records, while with first-order records an outcome can be supported if it has a corresponding unprofitable punishment. In the prisoner's dilemma, cooperation can be supported if and only if stage-game payoffs are either strictly supermodular and mild or strongly supermodular. The strength of the short-term coordination motive and the temptation to cheat thus determine the prospects for long-term cooperation.

We conclude by discussing some possible extensions and alternative models.

First-order records beyond the PD. Characterizing limit-supportable actions with first-order information in the absence of unprofitable punishments is a challenging problem. We solved this problem for the special case of cooperation in the PD, under equilibrium strictness and coordination-proofness. In an earlier version of this paper (Clark, Fudenberg, and Wolitzky, 2019a) we solved this problem for general stage games under a restriction to *trigger strategies*, where records are partitioned into two classes, one of which is absorbing. We found that such strategies can limit-support the play of an action a if and only if there exists a punishing action b that satisfies a generalized version of the definition of being an unprofitable punishment, where the requirement that $u(b, a) > u(a, a)$ is relaxed to $u(b, a) - u(a, a) < \min\{u(b, c) - u(a, c), u(a, a) - u(c, b)\}$. Extending this analysis beyond trigger strategies is a possible direction for future work, as is analyzing non-strict or non-coordination-proof strategies.

Simpler strategies. It is also interesting to consider simpler types of strategies. In Clark, Fudenberg, and Wolitzky (2020), we analyzed the performance of "tolerant"

grim trigger strategies without threading in the PD.¹⁹ We found that when $g < l/(1+l)$ such strategies can limit-support any cooperation share between g and $l/(1+l)$, and that otherwise they cannot limit-support any positive cooperation share.

Also, in earlier work (Clark, Fudenberg, and Wolitzky (2019a,b)) we considered strategies that condition only on the number of times a player was recorded as taking each action, and not the time sequence of these actions (as in the threading strategies used here). Such strategies yielded very similar but slightly more complicated results.

Sequential moves. In any strict equilibrium with first or second-order records, if players can “jump the gun” by taking their action before the opponent has a chance to respond, then only static equilibrium behavior can be supported.²⁰ However, our simultaneous-move specification applies not only when actions are literally simultaneous, but also whenever both players must choose their actions before fully observing their opponent’s action. This seems like a natural reduced-form model for the typical case where cooperation unfolds gradually within each match.²¹

Multiple populations. It is easy to adapt our model to settings with multiple populations of players. Here efficient outcomes can always be fully supported in situations with one-sided incentive problems.²² For example, suppose a population of player 1’s and a population of player 2’s repeatedly play a product choice game, where only player 1 faces binding moral hazard at the efficient action profile (and player 2 wants to match player 1’s action). The efficient outcome can always be supported with the following trigger strategies (with K chosen appropriately as a function of γ and ε): in each match, both partners play C if player 1’s score is less than K , and both play D if player 1’s score is greater than K .

¹⁹Tolerant grim strategies (Fudenberg, Rand, and Dreber (2012)) wait to punish until the opponent has defected several times.

²⁰To see why, note that by jumping the gun a player can obtain a stage-game payoff of $\max_r u(s(r, r'), s(r', r))$ when matched with an opponent with record r' , by taking action $s(\arg \max_r u(s(r, r'), s(r', r)), r')$. This implies that all players must receive the same payoff when matched with any given opponent.

²¹Also, if records are interdependent rather than second order, the strategies used to prove Theorem 2 remain equilibria for any possible move order.

²²Proposition 4 of Kandori (1992) is a similar result in a fixed-population model without noise.

Endogenous record systems. This paper has considered how features of an exogenously given record system determine the range of equilibrium outcomes. A natural next step is to endogenize the record system, for example by letting players strategically report their observations, either to a central database or directly to other individual players. Intuitively, first-order information is relatively easy to extract, since if a player is asked to report only their partner’s behavior, they have no reason to lie as this information does not affect their own future record. Whether and how society can obtain higher-order information is an interesting question for future study.²³

References

- [1] Charalambos D. Aliprantis and Kim C. Border. *Infinite Dimensional Analysis: A Hitchhiker’s Guide*. Third. Springer, 2006.
- [2] Nicolas Baumard. “Has Punishment Played a Role in the Evolution of Cooperation? A Critical Review”. *Mind & Society* 9 (2010), pp. 171–192.
- [3] V Bhaskar and Caroline Thomas. “Community Enforcement of Trust with Bounded Memory”. *Review of Economic Studies* 86 (2019), pp. 1010–1032.
- [4] Samuel Bowles and Herbert Gintis. *A Cooperative Species: Human Reciprocity and its Evolution*. Princeton University Press, 2011.
- [5] Robert Boyd et al. “The Evolution of Altruistic Punishment”. *Proceedings of the National Academy of Sciences* 100 (2003), pp. 3531–3535.
- [6] Luís Cabral and Ali Hortasçu. “The Dynamics of Seller Reputation: Evidence from eBay”. *The Journal of Industrial Economics* 58 (2010), pp. 54–78.
- [7] Daniel Clark, Drew Fudenberg, and Alexander Wolitzky. “Indirect Reciprocity with Simple Records”. *Proceedings of the National Academy of Sciences* 117 (2020), pp. 11344–11349.
- [8] Daniel Clark, Drew Fudenberg, and Alexander Wolitzky. “Steady-State Equilibria in Anonymous Repeated Games, I: Trigger Strategies in General Stage Games” (2019).

²³Other extensions include analyzing “optimal” record systems subject to some constraints, and the robustness of record systems to various manipulations, such as forging records of past interactions (as in Clark, Fudenberg, and Wolitzky (2019b)) or wiping one’s record clean by re-entering the game under a pseudonym (as in Friedman and Resnick (2001)).

- [9] Daniel Clark, Drew Fudenberg, and Alexander Wolitzky. “Steady-State Equilibria in Anonymous Repeated Games, II: Coordination-Proof Strategies in the Prisoner’s Dilemma” (2019).
- [10] Joyee Deb, Takuo Sugaya, and Alexander Wolitzky. “The Folk Theorem in Repeated Games with Anonymous Random Matching”. *Econometrica* 88 (2020), pp. 917–964.
- [11] Chrysanthos Dellarocas. “Reputation Mechanism Design in Online Trading Environments with Pure Moral Hazard”. *Information Systems Research* 16 (2005), pp. 209–230.
- [12] Francesc Dilmé. “Helping Behavior in Large Societies”. *International Economic Review* 57 (2016), pp. 1261–1278.
- [13] Glenn Ellison. “Cooperation in the Prisoner’s Dilemma with Anonymous Random Matching”. *Review of Economic Studies* 61 (1994), pp. 567–588.
- [14] Joseph Farrell and Eric Maskin. “Renegotiation in Repeated Games”. *Games and Economic Behavior* 1 (1989), pp. 327–360.
- [15] Eric J Friedman and Paul Resnick. “The Social Cost of Cheap Pseudonyms”. *Journal of Economics & Management Strategy* 10 (2001), pp. 173–199.
- [16] Drew Fudenberg and Kevin He. “Learning and Type Compatibility in Signaling Games”. *Econometrica* 86 (2018), pp. 1215–1255.
- [17] Drew Fudenberg, David Levine, and Eric Maskin. “The Folk Theorem with Imperfect Public Information”. *Econometrica* 62 (1994), pp. 997–1039.
- [18] Drew Fudenberg, David G. Rand, and Anna Dreber. “Slow to Anger and Fast to Forgive: Cooperation in an Uncertain World”. *American Economic Review* 102 (2012), pp. 720–749.
- [19] Avner Greif. “Contract Enforceability and Economic Institutions in Early Trade: The Maghribi Traders’ Coalition”. *The American Economic Review* 83 (1993), pp. 525–548.
- [20] Francesco Guala. “Reciprocity: Weak or Strong? What Punishment Experiments do (and do not) Demonstrate”. *Behavioral and Brain Sciences* 35 (2012), pp. 1–15.
- [21] Yuval Heller and Erik Mohlin. “Observations on Cooperation”. *Review of Economic Studies* 85 (2018), pp. 2253–2282.
- [22] Michihiro Kandori. “Social Norms and Community Enforcement”. *The Review of Economic Studies* 59 (1992), pp. 63–80.
- [23] Nobuhiro Kiyotaki and Randall Wright. “A Search-theoretic Approach to Monetary Economics”. *The American Economic Review* 83 (1993), pp. 63–77.

- [24] Benjamin Klein and Keith B Leffler. “The Role of Market Forces in Assuring Contractual Performance”. *Journal of Political Economy* 89 (1981), pp. 615–641.
- [25] Daniel B Klein. “Promise Keeping in the Great Society: A Model of Credit Information Sharing”. *Economics & Politics* 4 (1992), pp. 117–136.
- [26] Narayana R. Kocherlakota. “Money Is Memory”. *Journal of Economic Theory* 81 (1998), pp. 232–251.
- [27] Michael Luca. “Reviews, Reputation, and Revenue: The Case of Yelp.com” (2016).
- [28] Paul R 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 & Politics* 2 (1990), pp. 1–23.
- [29] Martin A. Nowak and Karl Sigmund. “Evolution of Indirect Reciprocity by Image Scoring”. *Nature* 393 (1998), pp. 573–577.
- [30] Masahiro Okuno-Fujiwara and Andrew Postlewaite. “Social Norms and Random Matching Games”. *Games and Economic Behavior* 9 (1995), pp. 79–109.
- [31] Elinor Ostrom. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge university press, 1990.
- [32] A. Jorge Padilla and Marco Pagano. “Sharing Default Information as a Borrower Discipline Device”. *European Economic Review* 44 (2000), pp. 1951–1980.
- [33] Karthik Panchanathan and Robert Boyd. “A Tale of Two Defectors: The Importance of Standing for Evolution of Indirect Reciprocity”. *Journal of Theoretical Biology* 224 (2003), pp. 115–126.
- [34] Paul Resnick and Richard Zeckhauser. “Trust among Strangers in Internet Transactions: Empirical Analysis of eBay’s Reputation System”. *The Economics of the Internet and E-commerce*. Emerald Group Publishing Limited, 2002, pp. 127–157.
- [35] Robert W. Rosenthal. “Sequences of Games with Varying Opponents”. *Econometrica* 47 (1979), pp. 1353–1366.
- [36] Robert W. Rosenthal and Henry J. Landau. “A Game-theoretic Analysis of Bargaining with Reputations”. *Journal of Mathematical Psychology* 20 (1979), pp. 233–255.
- [37] Robert Sugden. “New Developments in the Theory of Choice Under Uncertainty”. *Bulletin of Economic Research* 38 (1986), pp. 1–24.
- [38] Steven Tadelis. “Reputation and Feedback Systems in Online Platform Markets”. *Annual Review of Economics* 8 (2016), pp. 21–340.
- [39] Satoru Takahashi. “Community Enforcement when Players Observe Partners’ Past Play”. *Journal of Economic Theory* 145 (2010), pp. 42–62.

Appendix

A.1 Proof of Theorem 1

We first prove Theorem 1(i).

Without loss, relabel records so that two players with different ages can never share the same record. Let $R(t)$ be the set of feasible records for an age- t player, and fix a pairwise-public strategy σ . The proof relies on the following lemma.

Lemma 4. *If records are finite-partitional, there exists a family of finite subsets of R , $\{L(t, \eta)\}_{t \in \mathbb{N}, \eta > 0}$, such that*

1. $L(t, \eta) \subset R(t)$ for all $t \in \mathbb{N}, \eta > 0$,
2. For any $\mu \in \Delta(R)$, $\sum_{r \in L(0, \eta)} f_\sigma(\mu)[r] \geq (1 - \eta)(1 - \gamma)$ for all $\eta > 0$, and
3. For any $\mu \in \Delta(R)$ and $t > 0$, if $\sum_{r \in L(t-1, \eta)} \mu_r \geq (1 - \eta)(1 - \gamma)\gamma^{t-1}$ for all $\eta > 0$, then $\sum_{r \in L(t, \eta)} f_\sigma(\mu)[r] \geq (1 - \eta)(1 - \gamma)\gamma^t$ for all $\eta > 0$.

Proof. We construct the $\{L(t, \eta)\}$ by iteratively defining subfamilies of subsets of R that satisfy the necessary properties. First, take $L(0, \eta) = \{0\}$ for all $\eta > 0$. Conditions 1 and 2 are satisfied since $R(0) = \{0\}$ and $f_\sigma(\mu)[0] = 1 - \gamma$ for every $\mu \in \Delta(R)$.

Fix some t and take the subfamily of subsets corresponding to $t - 1$, that is $\{L(t - 1, \eta)\}_{\eta > 0}$. For every $\eta > 0$, consider the set of records $L(t - 1, \eta/2)$. Let $\lambda \in (0, 1)$ be such that $\lambda \geq (1 - \eta)/(1 - \eta/2)$. For any record $r \in L(t - 1, \eta/2)$, opposing record class R_m , and action profile $(a, a') \in A^2$, we can identify a finite set of “successor records” $S(r, m, a, a')$ such that a record r player who plays a against an opponent in class R_m playing a' moves to a record in $S(r, m, a, a')$ with probability greater than λ , i.e. $\sum_{r'' \in S(r, m, a, a')} \rho(r, r', a, a')[r''] \geq \lambda$ for all $r' \in R_m$. Let $L(t, \eta) = \bigcup_{r \in L(t-1, \eta/2)} \bigcup_{m \in \{1, \dots, M(r)\}} \bigcup_{(a, a') \in A^2} S(r, m, a, a')$. Note that $L(t, \eta)$ is finite and does not depend on μ . By construction, the probability that a surviving player with record in $L(t - 1, \eta/2)$ has a next-period record in $L(t, \eta)$ exceeds λ . For

any $\mu \in \Delta(R)$, it then follows that $\sum_{r \in L(t-1, \eta/2)} \mu_r \geq (1 - \eta/2)(1 - \gamma)\gamma^{t-1}$ implies $\sum_{r \in L(t, \eta)} f_\sigma(\mu)[r] \geq \lambda(1 - \eta/2)(1 - \gamma)\gamma^{t-1} \geq (1 - \eta)(1 - \gamma)\gamma^t$. \blacksquare

The next corollary is an immediate consequence of Properties 2 and 3 of Lemma 4.

Corollary 5. *For every $\mu \in \Delta(R)$ and $\eta > 0$, we have $\sum_{r \in L(t, \eta)} f_\sigma^{t'}(\mu)[r] \geq (1 - \eta)(1 - \gamma)\gamma^t$ for all $t' > t$, where $f_\sigma^{t'}$ denotes the t' th iterate of the update map f_σ .*

Fix a family $\{L(t, \eta)\}_{t \in \mathbb{N}, \eta > 0}$, satisfying the three properties in Lemma 4 and define \bar{M} , a subset of $\Delta(R)$, by

$$\bar{M} = \left\{ \mu \in \Delta(R) : \sum_{r \in R(t)} \mu_r = (1 - \gamma)\gamma^t \text{ and } \sum_{r \in L(t, \eta)} \mu_r \geq (1 - \eta)(1 - \gamma)\gamma^t \forall t \in \mathbb{N}, \eta > 0 \right\}.$$

Note that \bar{M} is convex and, by Corollary 5, must contain every steady-state distribution μ . The next lemma uses Corollary 5 to show that \bar{M} is non-empty.

Lemma 5. *There exists $\mu \in \Delta(R)$ satisfying $\sum_{r \in R(t)} \mu_r = (1 - \gamma)\gamma^t$ and $\sum_{r \in L(t, \eta)} \mu_r \geq (1 - \eta)(1 - \gamma)\gamma^t$ for every $t \in \mathbb{N}, \eta > 0$.*

Proof. Consider an arbitrary $\mu \in \Delta(R)$. Set $\mu^0 = \mu$, and, for every non-zero $i \in \mathbb{N}$, set $\mu^i = f_s(\mu^{i-1})$. Since R is countable, a standard diagonalization argument implies that there exists some $\tilde{\mu} \in [0, 1]^R$ and some subsequence $\{\mu^{i_j}\}_{j \in \mathbb{N}}$ such that $\lim_{j \rightarrow \infty} \mu_r^{i_j} = \tilde{\mu}_r$ for all $r \in R$.

For a given $t \in \mathbb{N}$, Corollary 5 implies that $\sum_{r \in L(t, \eta)} \mu_r^{i_j} \geq (1 - \eta)(1 - \gamma)\gamma^t$ for all $\eta > 0$ and all sufficiently high $j \in \mathbb{N}$, so

$$\sum_{r \in L(t, \eta)} \tilde{\mu}_r \geq (1 - \eta)(1 - \gamma)\gamma^t. \quad (1)$$

Moreover, for each $t \in \mathbb{N}$, $\sum_{r \in R(t)} \mu_r^{i_j} = (1 - \gamma)\gamma^t$ for all $j \in \mathbb{N}$, so $\sum_{r \in R(t)} \tilde{\mu}_r \leq (1 - \gamma)\gamma^t$. Since (1) holds for all $\eta \in (0, 1)$, this implies that $\sum_{r \in R(t)} \tilde{\mu}_r = (1 - \gamma)\gamma^t$, which together with (1) implies that $\tilde{\mu} \in \bar{M}$. \blacksquare

The following three claims imply that f_σ has a fixed point in \bar{M} ,²⁴ which completes the proof of Theorem 1(i).

Claim 1. \bar{M} is compact in the sup norm topology.

Claim 2. f_σ maps \bar{M} to itself.

Claim 3. f_σ is continuous in the sup norm topology.

Proof of Claim 1. Since \bar{M} is a metric space under the sup norm topology, it suffices to show that \bar{M} is sequentially compact. Consider a sequence $\{\mu^i\}_{i \in \mathbb{N}}$ of $\mu^i \in \bar{M}$. A similar argument to the proof of Lemma 5 shows that there exists some $\tilde{\mu} \in \bar{M}$ and some subsequence $\{\mu^{i_j}\}_{j \in \mathbb{N}}$ such that $\lim_{j \rightarrow \infty} \mu_r^{i_j} = \tilde{\mu}_r$ for all $r \in R$.

Here we show that $\lim_{j \rightarrow \infty} \mu^{i_j} = \tilde{\mu}$. For a given $\eta > 0$, there is a finite subset of records $L(\eta/2) \subset R$ such that $\sum_{r \in L(\eta/2)} \mu_r > 1 - \eta/2$ for every $\mu \in \bar{M}$. Thus, $|\mu_r^{i_j} - \tilde{\mu}_r| < \eta$ for all $r \notin L(\eta/2)$ for all $j \in \mathbb{N}$. Now let $J \in \mathbb{N}$ be such that $|\mu_r^{i_j} - \tilde{\mu}_r| < \eta$ for all $r \in L(\eta/2)$ whenever $j > J$. Then $\sup_{r \in R} |\mu_r^{i_j} - \tilde{\mu}_r| < \eta$ for all $j > J$. \blacksquare

Proof of Claim 2. For any $\mu \in \bar{M}$, Properties 2 and 3 of Lemma 4 imply that $\sum_{r \in L(t, \eta)} f_\sigma(\mu)[r] \geq (1 - \eta)(1 - \gamma)\gamma^t$ for all $t \in \mathbb{N}, \eta > 0$. Furthermore, $f_\sigma(\mu)[0] = 1 - \gamma$, and for all $t > 0$, $\gamma \sum_{r \in R(t-1)} \mu_r = \sum_{r \in R(t)} f_\sigma(\mu)[r]$, so $\sum_{r \in R(t-1)} \mu_r = (1 - \gamma)\gamma^{t-1}$ gives $\sum_{r \in R(t)} f_\sigma(\mu)[r] = (1 - \gamma)\gamma^t$. \blacksquare

Proof of Claim 3. Consider a sequence $\{\mu^i\}_{i \in \mathbb{N}}$ of $\mu^i \in \bar{M}$ with $\lim_{i \rightarrow \infty} \mu^i = \tilde{\mu} \in \bar{M}$. We will show that $\lim_{i \rightarrow \infty} f_\sigma(\mu^i) = f_\sigma(\tilde{\mu})$.

For any $\eta > 0$, there is a finite subset of records $L(\eta/4) \subset R$ such that $\sum_{r \in L(\eta/4)} \mu_r > 1 - \eta/4$ for every $\mu \in \bar{M}$. By Claim 2, $f_\sigma(\mu) \in \bar{M}$ for every $\mu \in \bar{M}$. The combination of these facts means that it suffices to show that $\lim_{i \rightarrow \infty} f_\sigma(\mu^i)[r] = f_\sigma(\tilde{\mu})[r]$ for all $r \in R$ to establish $\lim_{i \rightarrow \infty} f_\sigma(\mu^i) = f_\sigma(\tilde{\mu})$. Additionally, since $f_\sigma(\mu)[0] = 1 - \gamma$ is constant across $\mu \in \Delta(R)$, we need only consider the case where $r \neq 0$.

²⁴This follows from Corollary 17.56 (page 583) of Aliprantis and Border (2006), and noting that every normed space is a locally convex Hausdorff space.

For this case,

$$f_\sigma(\mu^i)[r] = \gamma \sum_{(r', r'') \in R^2} \mu_{r'}^i \mu_{r''}^i \phi(r', r'')[r],$$

and

$$f_\sigma(\tilde{\mu})[r] = \gamma \sum_{(r', r'') \in R^2} \tilde{\mu}_{r'} \tilde{\mu}_{r''} \phi(r', r'')[r].$$

Because $\sum_{r \in L(\eta/4)} \mu_r > 1 - \eta/4$ for every $\mu \in \bar{M}$, $\gamma \in (0, 1)$, and $0 \leq \phi(r', r'')[r] \leq 1$ for all $r', r'' \in R$, it follows that

$$\begin{aligned} |f_\sigma(\mu^i)[r] - f_\sigma(\tilde{\mu})[r]| &\leq \gamma \left| \sum_{(r', r'') \in L(\eta/4)^2} (\mu_{r'}^i \mu_{r''}^i - \tilde{\mu}_{r'} \tilde{\mu}_{r''}) \phi(r', r'')[r] \right| \\ &\quad + \gamma \left| \sum_{(r', r'') \notin L(\eta/4)^2} (\mu_{r'}^i \mu_{r''}^i - \tilde{\mu}_{r'} \tilde{\mu}_{r''}) \phi(r', r'')[r] \right| \\ &< \sum_{(r', r'') \in L(\eta/4)^2} |\mu_{r'}^i \mu_{r''}^i - \tilde{\mu}_{r'} \tilde{\mu}_{r''}| + \frac{1}{2}\eta. \end{aligned}$$

Since $\lim_{i \rightarrow \infty} \mu^i = \tilde{\mu}$, there exists some $I \in \mathbb{N}$ such that $\sum_{(r', r'') \in L(\eta/4)^2} |\mu_{r'}^i \mu_{r''}^i - \tilde{\mu}_{r'} \tilde{\mu}_{r''}| < \eta/2$ for all $i > I$, which gives $|f_\sigma(\mu^i)[r] - f_\sigma(\tilde{\mu})[r]| < \eta$ for all $i > I$. We thus conclude that $\lim_{i \rightarrow \infty} f_\sigma(\mu^i)[r] = f_\sigma(\tilde{\mu})[r]$. \blacksquare

We now prove Theorem 1(ii) by showing that no steady state exists when $\gamma > 1/2$ for the interdependent record system with $R = \mathbb{N}$ and $\rho(r, r', a, a') = \max\{r, r'\} + 1$. To see this, suppose toward a contradiction that μ is a steady state. Let r^* be the smallest record r such that $\sum_{r'=r}^\infty \mu_{r'} < 2 - 1/\gamma$, and let $\mu_* = \sum_{r=r^*}^\infty \mu_r < 2 - 1/\gamma$. Note that $\mu_* > 0$, as a player's record is no less than their age, so for any record threshold there is a positive measure of players whose records exceed the threshold.

Note that every surviving player with record $r \geq r^*$ retains a record higher than r^* , and at least fraction μ_* of the surviving players with record $r < r^*$ obtain a record higher than r^* (since this is the fraction of players with record $r < r^*$ that match with

a player with record $r \geq r^*$). Hence,

$$\sum_{r=r^*}^{\infty} f(\mu)[r] \geq \gamma\mu_* + \gamma(1 - \mu_*)\mu_* > \mu_*,$$

where the second inequality comes from $0 < \mu_* < 2 - 1/\gamma$. But in a steady state, $\sum_{r=r^*}^{\infty} f(\mu)[r] = \mu_*$, a contradiction.

A.2 Proof of Theorem 2

A.2.1 Proof of Theorem 2(i)

Let M be a positive integer such that $(u(a, a) - u(b, b))M > \max_{a'}\{\max\{u(a', a) - u(a, a), u(a', b) - u(b, b)\}\}$. Consider the following strategy, which we denote by σ^* .²⁵ Players begin in good standing. A player in good standing plays a when matched with fellow good-standing players and plays b against bad-standing players, while a player in bad standing always plays b . When a good-standing player's outcome is recorded as anything other than (a, a) or (b, b) , the player enters bad standing. A player remains in bad standing until they accumulate M (b, b) profiles, at which point they return to good standing. Claim 4 below shows that, under σ^* , the steady-state population share of a converges to 1 as $(\gamma, \varepsilon) \rightarrow (1, 0)$, and Claim 5 shows that σ^* induces strict equilibria when γ and ε are sufficiently close to their limit values.

Let $\tilde{\varepsilon}_{(a,a)} = \sum_{(\tilde{a}, \tilde{a}') \neq (a,a), (b,b)} \varepsilon_{(a,a), (\tilde{a}, \tilde{a}'})$ be the probability that the stage-game outcome is recorded as something other than (a, a) or (b, b) when the actual outcome is (a, a) , $\tilde{\varepsilon}_{(b,b)} = \sum_{(\tilde{a}, \tilde{a}') \neq (a,a), (b,b)} \varepsilon_{(b,b), (\tilde{a}, \tilde{a}'})$ be the probability that the outcome is recorded as something other than (a, a) or (b, b) when the actual outcome is (b, b) , and $\hat{\varepsilon}_{(b,b)} = \sum_{(\tilde{a}, \tilde{a}') \neq (b,b)} \varepsilon_{(b,b), (\tilde{a}, \tilde{a}'})$ be the probability that the outcome is recorded as something other than (b, b) when the actual outcome is (b, b) . In a steady state $\mu(\gamma, \varepsilon)$ for parameters (γ, ε) , let $\mu^G(\gamma, \varepsilon)$ be the corresponding share of good-standing players, and,

²⁵The strategy σ^* does not depend on (γ, ε) .

for $i \in \{0, \dots, M-1\}$, let $\mu^{B_i}(\gamma, \varepsilon)$ be the share of bad-standing players who have accumulated i (b, b) profiles since last entering bad standing.

Claim 4. *The unique limit point as $(\gamma, \varepsilon) \rightarrow (1, 0)$ of any sequence of steady-state shares induced by σ^* is $(\tilde{\mu}^G, \tilde{\mu}^{B_0}, \dots, \tilde{\mu}^{B_{M-1}}) = (1, 0, \dots, 0)$, so the share of good-standing players and the population share of action a converge to 1 in the $(\gamma, \varepsilon) \rightarrow (1, 0)$ limit.*

Proof. Let $(\tilde{\mu}^G, \tilde{\mu}^{B_0}, \dots, \tilde{\mu}^{B_{M-1}})$ be a limit point of a sequence of steady-state shares as $(\gamma, \varepsilon) \rightarrow (1, 0)$. The inflow into B_0 , the first phase of bad-standing, is $\gamma(\tilde{\varepsilon}_{a,a}\mu^G(\gamma, \varepsilon) + \tilde{\varepsilon}_{b,b}(1 - \mu^G(\gamma, \varepsilon)))\mu^G(\gamma, \varepsilon)$, which is the share of good-standing players that move into bad-standing in a given period. The outflow from B_0 is the sum of $(1 - \gamma)\mu^{B_0}(\gamma, \varepsilon)$, the share of players in phase B_0 who die in a given period, and $\gamma(1 - \hat{\varepsilon}_{(b,b)})\mu^{B_0}(\gamma, \varepsilon)$, the share of players in phase B_0 who move into phase B_1 in a given period. Thus, in a steady state, $\gamma(\tilde{\varepsilon}_{a,a}\mu^G(\gamma, \varepsilon) + \tilde{\varepsilon}_{b,b}(1 - \mu^G(\gamma, \varepsilon)))\mu^G(\gamma, \varepsilon) = (1 - \gamma\hat{\varepsilon}_{(b,b)})\mu^{B_0}(\gamma, \varepsilon)$. Taking the limit of this equation as $(\gamma, \varepsilon) \rightarrow (1, 0)$ gives $\tilde{\mu}^{B_0} = 0$. Likewise, equating the inflow and outflows of phase B_i for $0 < i < M$ gives $\gamma(1 - \hat{\varepsilon}_{(b,b)})\mu^{B_{i-1}}(\gamma, \varepsilon) = (1 - \gamma\hat{\varepsilon}_{(b,b)})\mu^{B_i}(\gamma, \varepsilon)$, and taking the limit of this equation as $(\gamma, \varepsilon) \rightarrow (1, 0)$ shows that $\tilde{\mu}^{B_i} = \tilde{\mu}^{B_{i-1}}$. Combining this with $\tilde{\mu}^{B_0} = 0$ gives $\tilde{\mu}^{B_i} = 0$ for all $i \in \{0, \dots, M-1\}$. Since the good-standing population share and bad-standing population shares always sum to 1, it follows that $\tilde{\mu}^G = 1$. ■

Claim 5. *Given any sequence of steady states $\mu(\gamma, \varepsilon)$ induced by σ^* , $(\sigma^*, \mu(\gamma, \varepsilon))$ is a strict equilibrium when γ is sufficiently close to 1 and ε is sufficiently close to 0.*

Proof. For $0 \leq i < M-1$, the value functions in the bad-standing phase B_i and the subsequent bad-standing phase B_{i+1} satisfy

$$V^{B_i} = (1 - \gamma)u(b, b) + \gamma\hat{\varepsilon}_{(b,b)}V^{B_i} + \gamma(1 - \hat{\varepsilon}_{(b,b)})V^{B_{i+1}}. \quad (2)$$

Similarly the value functions in the final bad-standing phase B_{M-1} and the good-

standing phase G are linked by

$$V^{B_{M-1}} = (1 - \gamma)u(b, b) + \gamma\hat{\varepsilon}_{(b,b)}V^{B_{M-1}} + \gamma(1 - \hat{\varepsilon}_{(b,b)})V^G. \quad (3)$$

The value function in the good-standing phase equals the average flow payoff in the population in a given period (since newborn players are in good standing), so $V^G = \mu^G(\gamma, \varepsilon)^2u(a, a) + (1 - \mu^G(\gamma, \varepsilon)^2)u(b, b)$. Combining this with $\lim_{(\gamma, \varepsilon) \rightarrow (1, 0)} \mu^G(\gamma, \varepsilon) = 1$, we obtain $\lim_{(\gamma, \varepsilon) \rightarrow (1, 0)} V^G = u(a, a)$. Taking the limits of Equations 2 and 3 as $(\gamma, \varepsilon) \rightarrow (1, 0)$ thus gives $\lim_{(\gamma, \varepsilon) \rightarrow (1, 0)} V^{B_i} = \lim_{(\gamma, \varepsilon) \rightarrow (1, 0)} V^G = u(a, a)$ for all $i \in \{0, \dots, M - 1\}$.

A player in bad-standing phase i where $0 \leq i < M - 1$ strictly prefers to play b against b when $(1 - \gamma)u(b, b) + \gamma\hat{\varepsilon}_{(b,b)}V^{B_i} + \gamma(1 - \hat{\varepsilon}_{(b,b)})V^{B_{i+1}} > (1 - \gamma)u(a', b) + \gamma(1 - \varepsilon_{(a',b),(b,b)})V^{B_i} + \gamma\varepsilon_{(a',b),(b,b)}V^{B_{i+1}}$ holds for $a' \neq b$. Manipulating this gives $(1 - \hat{\varepsilon}_{(b,b)} - \varepsilon_{(a',b),(b,b)})\gamma(V^{B_{i+1}} - V^{B_i})/(1 - \gamma) > u(a', b) - u(b, b)$. Equation 2 can be rewritten as

$$\frac{\gamma}{1 - \gamma}(V^{B_{i+1}} - V^{B_i}) = \frac{\gamma}{1 - \gamma\hat{\varepsilon}_{(b,b)}}(V^{B_{i+1}} - u(b, b)),$$

so we obtain $\lim_{(\gamma, \varepsilon) \rightarrow (1, 0)} (1 - \hat{\varepsilon}_{(b,b)} - \varepsilon_{(a',b),(b,b)})\gamma(V^{B_{i+1}} - V^{B_i})/(1 - \gamma) = u(a, a) - u(b, b)$. Since $\max_{a'} u(a', b) < u(a, a)$, it follows that the incentives of players in bad-standing phase i are satisfied for (γ, ε) sufficiently close to $(1, 0)$.

An almost identical argument shows that the incentives of players in bad-standing phase $M - 1$ are satisfied for (γ, ε) sufficiently close to $(1, 0)$. Thus, all that remains is to show that the incentives of players in good-standing are satisfied in the limit. A good-standing player has strict incentives to play a against a when $(1 - \gamma)u(a, a) + \gamma(1 - \tilde{\varepsilon}_{(a,a)})V^G + \gamma\tilde{\varepsilon}_{(a,a)}V^{B_0} > (1 - \gamma)u(a', a) + \gamma(\varepsilon_{(a',a),(a,a)} + \varepsilon_{(a',a),(b,b)})V^G + \gamma(1 - \varepsilon_{(a',a),(a,a)} - \varepsilon_{(a',a),(b,b)})V^{B_0}$ holds for $a' \neq a$. Manipulating this gives $(1 - \tilde{\varepsilon}_{(a,a)} - \varepsilon_{(a',a),(a,a)} - \varepsilon_{(a',a),(b,b)})\gamma(V^G - V^{B_0})/(1 - \gamma) > u(a', a) - u(a, a)$. Similarly, a good-standing player has strict incentives to play b against b when $(1 - \gamma)u(b, b) + \gamma(1 - \tilde{\varepsilon}_{(b,b)})V^G + \gamma\tilde{\varepsilon}_{(b,b)}V^{B_0} > (1 - \gamma)u(a', b) + \gamma(\varepsilon_{(a',b),(a,a)} + \varepsilon_{(a',b),(b,b)})V^G + \gamma(1 - \varepsilon_{(a',b),(a,a)} - \varepsilon_{(a',b),(b,b)})V^{B_0}$ holds for $a' \neq b$. Manipulating this gives $(1 - \tilde{\varepsilon}_{(b,b)} - \varepsilon_{(a',a),(a,a)} - \varepsilon_{(a',a),(b,b)})\gamma(V^G - V^{B_0})/(1 - \gamma) >$

$u(a', b) - u(b, b)$. Combining Equations 2 and 3 gives

$$\frac{\gamma}{1-\gamma}(V^G - V^{B_0}) = \gamma \frac{1 - \left(\frac{\gamma(1-\hat{\varepsilon}(b,b))}{1-\gamma\hat{\varepsilon}(b,b)}\right)^M}{1-\gamma}(V^G - u(b, b)).$$

It follows that $\lim_{(\gamma, \varepsilon) \rightarrow (1, 0)} (1 - \tilde{\varepsilon}_{(a,a)} - \varepsilon_{(a',a),(a,a)} - \varepsilon_{(a',a),(b,b)})\gamma(V^G - V^{B_0})/(1-\gamma) = M(u(a, a) - u(b, b))$. Since $M(u(a, a) - u(b, b)) > \max_{a'} \{\max\{u(a', a) - u(a, a), u(a', b) - u(b, b)\}\}$, good-standing players' incentives are satisfied for (γ, ε) sufficiently close to $(1, 0)$. \blacksquare

A.2.2 Proof of Theorem 2(ii)

We first show that a can be limit-supported and then show that the constructed equilibria are coordination-proof.

A.2.2.1 Proof that a is Limit-Supported by Strict Equilibria

Let $0 < \underline{\gamma} < \bar{\gamma} < 1$ be such that

$$\frac{\gamma}{1-\gamma} > \max_{(x_1, x_2)} \frac{u(x_1, x_2) - u(b, b)}{u(a, a) - u(b, b)} \quad (4)$$

for all $\gamma \in [\underline{\gamma}, \bar{\gamma}]$. Denote the grim trigger strategy described in Section 3 by σ_{tr} , and let μ^G denote the share of good-standing players in a steady state. We will show that for all $\delta > 0$, there is an $\bar{\varepsilon} > 0$ such that, whenever $\gamma \in [\underline{\gamma}, \bar{\gamma}]$ and $\varepsilon_{(x_1, x_2), (x'_1, x'_2)} < \bar{\varepsilon}$ for all $(x_1, x_2), (x'_1, x'_2) \in A^2$, σ_{tr} induces strict equilibria satisfying $\mu^G > 1 - \delta$. Thus, σ_{tr} can be combined with the threading technique described in the text to limit-support a as $(\gamma, \varepsilon) \rightarrow (1, 0)$.

Claim 6 below shows that the steady-state share of good-standing players induced by σ_{tr} converges to 1 uniformly over $\gamma \in [\underline{\gamma}, \bar{\gamma}]$ as $\varepsilon \rightarrow 0$. For the remainder of the proof, we restrict attention to $\gamma \in [\underline{\gamma}, \bar{\gamma}]$. Claim 7 then shows that the incentives of good-standing players are satisfied when ε is sufficiently small. These two claims

together complete the argument, as the incentives of bad-standing players are always satisfied since (b, b) is a strict static equilibrium.

Claim 6. *For all $\delta > 0$, there is an $\bar{\varepsilon} > 0$ such that, whenever $\varepsilon_{(x_1, x_2), (x'_1, x'_2)} < \bar{\varepsilon}$ for all $(x_1, x_2), (x'_1, x'_2) \in A^2$, the steady states induced by σ_{tr} satisfy $\mu^G > 1 - \delta$.*

Proof. Note that the inflow into good standing is simply $1 - \gamma$, the share of newborn players. The outflow from good standing is the sum of $(1 - \gamma)\mu^G$, the share of good-standing players who die in a given period, and $\gamma(\tilde{\varepsilon}_{(a,a)}\mu^G + \tilde{\varepsilon}_{(b,b)}(1 - \mu^G))\mu^G$, the share of good-standing players whose outcome is recorded as something other than (a, a) or (b, b) in a given period. In a steady state, these inflows and outflows must be equal, and setting the corresponding expressions equal to each other gives

$$\mu^G = \frac{1 - \gamma}{1 - \gamma + \gamma(\tilde{\varepsilon}_{(a,a)}\mu^G + \tilde{\varepsilon}_{(b,b)}(1 - \mu^G))} \geq \frac{1 - \gamma}{1 - \gamma + \gamma \max\{\tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}\}}.$$

The claim then follows since $\lim_{\varepsilon \rightarrow 0} \inf_{\gamma \in [\underline{\gamma}, \bar{\gamma}]} (1 - \gamma) / (1 - \gamma + \gamma \max\{\tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}\}) = 1$. ■

Claim 7. *For all $\delta > 0$, there is an $\bar{\varepsilon} > 0$ such that, whenever $\varepsilon_{(x_1, x_2), (x'_1, x'_2)} < \bar{\varepsilon}$ for all $(x_1, x_2), (x'_1, x'_2) \in A^2$, the incentives of good-standing players are satisfied.*

Proof. The value function of good-standing players, V^G , equals the average flow payoff in the population in a given period (since newborn players are in good standing), so $V^G = \mu^G(\mu^G u(a, a) + (1 - \mu^G)u(b, b)) + (1 - \mu^G)u(b, b)$. In contrast, the value function of bad-standing players, V^B , equals the expected flow payoff of bad-standing players, so $V^B = u(b, b)$.

When facing an opponent playing a , the expected payoff of a good-standing player from playing a is $(1 - \gamma)u(a, a) + \gamma(1 - \tilde{\varepsilon}_{(a,a)})V^G + \tilde{\varepsilon}_{(a,a)}V^B$ while their expected payoff from playing $x \neq a$ is $(1 - \gamma)u(x, a) + \gamma(\varepsilon_{(x,a), (a,a)} + \varepsilon_{(x,a), (b,b)})V^G + (1 - \varepsilon_{(x,a), (a,a)} - \varepsilon_{(x,a), (b,b)})V^B$. Thus, a good-standing player strictly prefers to play a rather than any $x \neq a$ precisely when

$$\frac{\gamma}{1 - \gamma} > \max_{x \neq a} \frac{u(x, a) - u(a, a)}{(1 - \tilde{\varepsilon}_{(a,a)} - \varepsilon_{(x,a), (a,a)} - \varepsilon_{(x,a), (b,b)})\mu^G(u(a, a) - u(b, b))}.$$

Claim 6 implies that, as $\varepsilon \rightarrow 0$, the right-hand side of this inequality converges uniformly to $\max_{x \neq a} (u(x, a) - u(a, a)) / (u(a, a) - u(b, b))$. By the inequality in (4), we conclude that a good-standing player strictly prefers to match a with a instead of playing some $x \neq a$ for sufficiently small noise.

When facing an opponent playing b , the expected payoff of a good-standing player from playing b is $(1 - \gamma)u(b, b) + \gamma(1 - \tilde{\varepsilon}_{(b,b)})V^G + \tilde{\varepsilon}_{(b,b)}V^B$ while their expected payoff from playing $x \neq b$ is $(1 - \gamma)u(x, b) + \gamma(\varepsilon_{(x,b),(a,a)} + \varepsilon_{(x,b),(b,b)})V^G + (1 - \varepsilon_{(x,b),(a,a)} - \varepsilon_{(x,b),(b,b)})V^B$. Thus a good-standing player strictly prefers to play a rather than any $x \neq b$ precisely when

$$\frac{\gamma}{1 - \gamma} > \max_{x \neq b} \frac{u(x, b) - u(b, b)}{(1 - \tilde{\varepsilon}_{(b,b)} - \varepsilon_{(x,b),(a,a)} - \varepsilon_{(x,b),(b,b)})\mu^G(u(a, a) - u(b, b))}.$$

Claim 6 implies that, as $\varepsilon \rightarrow 0$, the right-hand side of this inequality converges uniformly to $\max_{x \neq b} (u(x, b) - u(b, b)) / (u(a, a) - u(b, b))$. By the inequality in (4), we conclude that a good-standing player strictly prefers to match b with b instead of playing some $x \neq b$ for sufficiently small noise. \blacksquare

A.2.2.2 Proof of Coordination-Proofness

We first argue that in every match between bad-standing players, there is no Nash equilibrium in the augmented game that Pareto-dominates (b, b) . Note that the outcome of the current match does not affect a bad-standing player's continuation value. Thus, any Nash equilibrium in the augmented game between two bad-standing players must also be a static equilibrium in the stage game. Since there is no static equilibrium that Pareto-dominates (b, b) , it follows that two bad-standing players playing (b, b) is coordination-proof.

Now we show that in any match involving a good-standing player, there is no Nash equilibrium in the augmented game that Pareto-dominates the action profile the players are supposed to play. By an argument similar to the one showing that a good-standing player strictly prefers to play a against a , no good-standing player would ever prefer an

action profile other than (a, a) or (b, b) be played in one of their matches when $\gamma \in [\underline{\gamma}, \bar{\gamma}]$ and noise is sufficiently small. Thus, in any match involving a good-standing player, we need only consider whether (a, a) or (b, b) are Nash equilibria in the augmented game and whether one of these profiles Pareto-dominates the other. When two good-standing players match, both (a, a) and (b, b) are Nash equilibria in the augmented game, but (b, b) does not Pareto-dominate (a, a) . Indeed, if (b, b) did Pareto-dominate (a, a) , this would imply that the value functions for these good-standing players would be no higher than $u(b, b)$, which is not possible given that $u(a, a) > u(b, b)$. Thus, the prescribed play between two good-standing players is coordination-proof. Moreover, in any match involving a bad-standing player, all Nash equilibria in the augmented game require the bad-standing player to play a static best-response to the action of their opponent. Because $u(a, a) > u(b, b)$ and (b, b) is not Pareto-dominated by any static equilibrium, (a, a) is not a static equilibrium, so (b, b) is coordination-proof when a good-standing player matches a bad-standing player.

A.3 Proof of Theorem 5

Section A.3.1 derives the incentive constraints that must be satisfied in any strict equilibrium with noisy first-order records, and Section A.3.2 proves Theorem 5(i) (necessary conditions for cooperation). The main step is proving Lemma 9, which shows that $\mu^P + \mu^S(l - g) > g$ in any strict, coordination-proof equilibrium with $\mu^C > 0$. Section A.3.3 proves Theorem 5(ii) (sufficient conditions for cooperation). This part of the proof is split into three parts: Section A.3.3.1 shows that threaded grim trigger strategies can limit-support cooperation when $g < 1$; Section A.3.3.2 shows that threaded “defector→preciprocator→supercooperator→defector” strategies can limit-support cooperation when $l > g + g^2$; and Section A.3.3.3 shows that each of these classes of equilibria is coordination-proof.

A.3.1 Incentive Constraints with Noisy Records

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 , and $(D|D)_r$ denotes the condition that D 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 expectation 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 r player. Note that $p_r^D > 0$ for all r since $\varepsilon_C(r) > 0$ and $\varepsilon_D(r) < 1$.

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

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

Lemma 6. *For any noisy record system,*

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

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 $(1 - \varepsilon_C(r))V_r^C + \varepsilon_C(r)V_r^D$, since a recording of C is fed into the record system with probability $1 - \varepsilon_C(r)$ and a recording of D is fed into the record

system with probability $\varepsilon_C(r)$. Similarly, when the player plays D , their expected continuation payoff is $\varepsilon_D(r)V_r^C + (1 - \varepsilon_D(r))V_r^D$. Thus, the $(C|C)_r$ constraint is $1 - \gamma + \gamma(1 - \varepsilon_C(r))V_r^C + \gamma\varepsilon_C(r)V_r^D > (1 - \gamma)(1 + g) + \gamma\varepsilon_D(r)V_r^C + (1 - \varepsilon_D(r))V_r^D$, which is equivalent to

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

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. \blacksquare

The strategies we use to prove part (ii) of the theorem depend on a player's record only through their age and their "score," which is the number of times they have been recorded as playing D . For such *scoring strategies*, we slightly abuse notation in writing V_k for the continuation payoff of a player with score k .²⁶ The incentive constraints for such strategies take a relatively simple form with the canonical first-order record system: For all k we have $\varepsilon_C(k) = \varepsilon_C$, $\varepsilon_D(k) = \varepsilon_D$, $V_k^C = V_k$, and $V_k^D = V_{k+1}$, so the normalized reward simplifies to

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

Lemma 7. *For scoring strategies and the canonical first-order record system, Lemma 6 holds with $W_k = (1 - \varepsilon_C - \varepsilon_D)(\pi_k - V_k)/p_k^D$.*

A.3.2 Proof of Theorem 5(i)

Theorem 5(i) follows from the following two lemmas.

Lemma 8. *For any first-order record system, in any strict equilibrium, $\mu^S < 1/(1 + g)$.*

Lemma 9. *For any noisy first-order record system, in any strict, coordination-proof equilibrium with $\mu^C > 0$, $\mu^P + \mu^S(l - g) > g$.*

²⁶Recall that V_r is defined as the continuation value of a player with record r . Under scoring strategies, two players with different records that share the same score have the same continuation value, so we can index V by k rather than r .

Lemma 8 says that there cannot be too many supercooperators. It holds because new players with record 0 have the option of always playing D , so in any strict equilibrium with $\mu^C > 0$, it must be that $\mu^S(1+g) < V_0 \leq 1$, which gives $\mu^S < 1/(1+g)$.

Conversely, Lemma 9 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. It is proved in the next subsection.

Theorem 5(i) follows from Lemmas 8 and 9 since, by Lemma 3, it is impossible to satisfy both $\mu^S < 1/(1+g)$ and $\mu^P + \mu^S(l-g) > g$ when $g \geq 1$ and $l \leq g+g^2$.

A.3.2.1 Necessary Conditions for Cooperation and Proof of Lemma 9

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 10. *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 otherwise. 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 $\mu^C = 0$. ■

Lemma 11. *If $\mu^C > 0$, there is some record r' that is a preciprocator or a supercooperator and satisfies*

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

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 10, such an 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 10, for sufficiently high n , the record r_n must be either a preciprocator or a supercooperator. \blacksquare

Proof of Lemma 9. First, take the case where the record r' identified in Lemma 11 is a preciprocator. Then by Lemma 6, we must have

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

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

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

Note that $p_{r'}^D \geq (1 - \varepsilon_D(r'))\mu^D$ since a preciprocator 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 $1 - \varepsilon_D(r')$. This gives $(1 - \varepsilon_C(r') - \varepsilon_D(r'))\mu^D/p_{r'}^D < 1$, so $\mu^P + \mu^S(l-g) > g$ must hold.

Now take the case where r' is a supercooperator. By Lemma 6, $\pi_{r'} - V_{r'} + (\gamma/(1-\gamma))(V_{r'}^C - V_{r'}) > 0$. When $\pi_{r'} = \mu^C - \mu^D l$ and $V_{r'} - (\gamma/(1-\gamma))(V_{r'}^C - V_{r'}) \geq \mu^P \mu^C +$

$\mu^S(\mu^C - \mu^D l) + \mu^D \mu^S(1 + g)$, this implies that

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

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

A.3.3 Proof of Theorem 5(ii)

A.3.3.1 Limit-Supporting C when $g < 1$

Let $0 < \underline{\gamma} < \bar{\gamma} < 1/2$ be such that

$$g < \frac{\gamma}{1 - \gamma} < l \tag{5}$$

for all $\gamma \in [\underline{\gamma}, \bar{\gamma}]$. Consider the grim trigger strategy, and let $\mu^C(\gamma, \varepsilon)$ denote the steady state share of cooperators, i.e. those players who have not been recorded as playing D , for parameters (γ, ε) . (As we will see, there is a unique steady state when noise is sufficiently small.) We will show that for all $\delta > 0$, there is an $\bar{\varepsilon} > 0$ such that, whenever $\gamma \in [\underline{\gamma}, \bar{\gamma}]$ and $\varepsilon_C, \varepsilon_D < \bar{\varepsilon}$, grim trigger induces strict equilibria satisfying $\mu^C(\gamma, \varepsilon) > 1 - \delta$. Thus, grim trigger can be combined with threading to limit-support C as $(\gamma, \varepsilon) \rightarrow (1, 0)$.

First, we establish that for all $\delta > 0$, there is an $\bar{\varepsilon} > 0$ such that whenever $\gamma \in [\underline{\gamma}, \bar{\gamma}]$ and $\varepsilon_C, \varepsilon_D < \bar{\varepsilon}$, there is a unique steady state induced by grim trigger and it satisfies $\mu^C > 1 - \delta$. Note that the inflow into cooperator status is $1 - \gamma$, the share of newborn players. The outflow from cooperator status is the sum of $(1 - \gamma)\mu^C$, the share of cooperators who die in a given period, and $\gamma(\varepsilon_C \mu^C + (1 - \varepsilon_D)(1 - \mu^C))\mu^C$, the share of cooperators who are recorded as playing D in a given period. In a steady state, these inflows and outflows must be equal, so

$$1 - \gamma = (1 - \gamma + \gamma(\varepsilon_C \mu^C + (1 - \varepsilon_D)(1 - \mu^C)))\mu^C.$$

This expression has a unique solution $\mu^C \in [0, 1]$ when ε_C and ε_D are sufficiently small, given by

$$\mu^C(\gamma, \varepsilon) = \frac{1 - \gamma\varepsilon_D - \sqrt{(1 - \gamma\varepsilon_D)^2 - 4\gamma(1 - \varepsilon_C - \varepsilon_D)(1 - \gamma)}}{2\gamma(1 - \varepsilon_C - \varepsilon_D)}.$$

Note that $\mu^C(\gamma, \varepsilon)$ is continuous for $\gamma \in [\underline{\gamma}, \bar{\gamma}]$ and sufficiently small $\varepsilon_C, \varepsilon_D$, and $\mu^C(\gamma, 0) = 1$ for all $\gamma \leq 1/2$. It follows that there is an $\bar{\varepsilon} > 0$ such that $\mu^C(\gamma, \varepsilon) > 1 - \delta$ for all $\gamma \in [\underline{\gamma}, \bar{\gamma}]$ and $\varepsilon_C, \varepsilon_D < \bar{\varepsilon}$.

Now we establish that for all $\delta > 0$, there is an $\bar{\varepsilon} > 0$ such that whenever $\gamma \in [\underline{\gamma}, \bar{\gamma}]$ and $\varepsilon_C, \varepsilon_D < \bar{\varepsilon}$, the incentives of reciprocators are satisfied. (The incentives of defectors are clearly satisfied.) We will use the facts that the value function of reciprocators, V^C , equals the average flow payoff in the population in a given period, $(\mu^C(\gamma, \varepsilon))^2$, and that the value function of defectors is $V^D = 0$.

When facing an opponent playing C , the expected payoff for a reciprocator from playing C is $1 - \gamma + \gamma(1 - \varepsilon_C)(\mu^C(\gamma, \varepsilon))^2$ while their expected payoff from playing D is $(1 - \gamma)(1 + g) + \gamma\varepsilon_D(\mu^C(\gamma, \varepsilon))^2$. Thus, a reciprocator strictly prefers to play C against an opponent playing C if and only if $1 - \gamma + \gamma(1 - \varepsilon_C)(\mu^C(\gamma, \varepsilon))^2 > (1 - \gamma)(1 + g) + \gamma\varepsilon_D(\mu^C(\gamma, \varepsilon))^2$, which simplifies to

$$\frac{\gamma}{1 - \gamma} > \frac{g}{(1 - \varepsilon_C - \varepsilon_D)(\mu^C(\gamma, \varepsilon))^2}.$$

When facing an opponent playing D , the expected payoff to C for a reciprocator is $-(1 - \gamma)l + \gamma(1 - \varepsilon_C)(\mu^C(\gamma, \varepsilon))^2$ while their expected payoff from playing D is $\gamma\varepsilon_D(\mu^C(\gamma, \varepsilon))^2$. Thus, a reciprocator strictly prefers to play D against an opponent playing D if and only if $-(1 - \gamma)l + \gamma(1 - \varepsilon_C)(\mu^C(\gamma, \varepsilon))^2 < \gamma\varepsilon_D(\mu^C(\gamma, \varepsilon))^2$, which simplifies to

$$\frac{\gamma}{1 - \gamma} < \frac{l}{(1 - \varepsilon_C - \varepsilon_D)(\mu^C(\gamma, \varepsilon))^2}.$$

Combining these incentive conditions shows that all the incentives of a reciprocator

are satisfied if and only if

$$\frac{g}{(1 - \varepsilon_C - \varepsilon_D)(\mu^C(\gamma, \varepsilon))^2} < \frac{\gamma}{1 - \gamma} < \frac{l}{(1 - \varepsilon_C - \varepsilon_D)(\mu^C(\gamma, \varepsilon))^2}.$$

As $\varepsilon \rightarrow 0$, the left-most expression and right-most expression in this inequality converge uniformly to g and l , respectively. From (5), we conclude that the incentives of a preciprocator are satisfied for sufficiently small noise when $\gamma \in [\underline{\gamma}, \bar{\gamma}]$.

A.3.3.2 Limit-Supporting C when $l > g + g^2$

We use the class of strategies of the form “ $D_J P_K S_1 D_\infty$,” where $J, K \in \mathbb{N}$. These strategies specify that a player is a defector until they have been recorded as playing D J times. Subsequently, the player is a preciprocator until they have been recorded as playing D K more times, and then a supercooperator until they are recorded as playing D once more, after which they permanently become a defector. Throughout, we let μ^{D_1} be the share of players who have been recorded as playing D fewer than J times (and are thus defectors), μ^P be the share of preciprocators (those with score $J \leq k < J + K$), μ^S be the share of supercooperators (those with score $k = J + K$), and μ^{D_2} be the share of defectors with a score $k > J + K$. We also let $\mu^C = \mu^P + \mu^S$ be the total share of cooperators and $\mu^D = \mu^{D_1} + \mu^{D_2} = 1 - \mu^C$ be the total share of defectors. We will show that for all $\delta > 0$, there are $0 < \underline{\gamma} < \bar{\gamma} < 1$ and $\bar{\varepsilon} > 0$ such that when $\gamma \in [\underline{\gamma}, \bar{\gamma}]$ and $\varepsilon_C, \varepsilon_D < \bar{\varepsilon}$, this strategy class gives equilibria satisfying $\mu^C > 1 - \delta$. Thus, these strategies can be combined with threading to limit-support C as $(\gamma, \varepsilon) \rightarrow (1, 0)$.

The following lemma characterizes precisely which population shares and parameters are consistent with an equilibrium using a $D_J P_K S_1 D_\infty$ strategy. The statement of the lemma involves the functions $\alpha : (0, 1) \times (0, 1) \rightarrow (0, 1)$ and $\beta : (0, 1) \times (0, 1) \times$

$[0, 1] \rightarrow (0, 1)$, defined by

$$\alpha(\gamma, \psi) = \frac{\gamma\psi}{1 - \gamma + \gamma\psi},$$

$$\beta(\gamma, \varepsilon, \mu^D) = \frac{\gamma(\varepsilon_C + (1 - \varepsilon_C - \varepsilon_D)\mu^D)}{1 - \gamma + \gamma(\varepsilon_C + (1 - \varepsilon_C - \varepsilon_D)\mu^D)}.$$

Lemma 12. *There is a $D_J P_K S_1 D_\infty$ equilibrium with shares μ^{D_1} , μ^P , μ^S , and μ^{D_2} if and only if the following conditions hold:*

1. *Feasibility:*

$$\mu^{D_1} = 1 - \alpha(\gamma, 1 - \varepsilon_D)^J,$$

$$\mu^P = \alpha(\gamma, 1 - \varepsilon_D)^J (1 - \beta(\gamma, \varepsilon, \mu^D)^K),$$

$$\mu^S = \alpha(\gamma, 1 - \varepsilon_D)^J \beta(\gamma, \varepsilon, \mu^D)^K (1 - \alpha(\gamma, \varepsilon_C)),$$

$$\mu^{D_2} = \alpha(\gamma, 1 - \varepsilon_D)^J \beta(\gamma, \varepsilon, \mu^D)^K \alpha(\gamma, \varepsilon_C).$$
2. *Incentives:*

$$(C|C)_J : \frac{(1 - \varepsilon_C - \varepsilon_D)\mu^D}{\varepsilon_C + (1 - \varepsilon_C - \varepsilon_D)\mu^D} \left(\frac{\mu^S}{1 - \mu^{D_1}} l + \frac{\mu^{D_2}}{\mu^{D_1}(1 - \mu^{D_1})} (\mu^P - \mu^S g) \right) > g,$$

$$(D|D)_{J+K-1} : \frac{\gamma(1 - \varepsilon_C - \varepsilon_D)((1 - \alpha(\gamma, \varepsilon_C))\mu^D l + \alpha(\gamma, \varepsilon_C)(\mu^P - \mu^S g))}{1 - \gamma + \gamma(\varepsilon_C + (1 - \varepsilon_C - \varepsilon_D)\mu^D)} < l,$$

$$(C|D)_{J+K} \text{ (if } \mu^S > 0\text{)} : \frac{\gamma(1 - \varepsilon_C - \varepsilon_D)}{1 - \gamma + \gamma\varepsilon_C} (\mu^P - \mu^S g - \mu^D l) > l.$$

The proof of Lemma 12 is in OA.3.1. The feasibility constraints come from calculating the relevant steady-state shares for the strategy $D_J P_K S_1 D_\infty$. The $(C|C)_J$ incentive constraint comes from solving V_J and using Lemma 7. The $(D|D)_{J+K-1}$ and $(C|D)_{J+K}$ constraints are derived by relating the value functions of adjacent records.

Since $l > g + g^2$, it can be shown that, for all sufficiently small $\delta > 0$, there are $\bar{\mu}^P, \bar{\mu}^S > 0$ satisfying $\bar{\mu}^P + \bar{\mu}^S = 1 - \delta$, $\bar{\mu}^S > (g/l)(1 - \delta)$, and $\bar{\mu}^P - \bar{\mu}^S g - \delta l > 0$. Fix such a δ and the corresponding $\bar{\mu}^P, \bar{\mu}^S$. There is some sufficiently small $\eta \in (0, \delta/2)$ such that the above inequalities hold when $\bar{\mu}^P, \bar{\mu}^S$, and δ are respectively replaced with any μ^P, μ^S , and $\tilde{\delta}$ satisfying $|\mu^P - \bar{\mu}^P|, |\mu^S - \bar{\mu}^S| \leq \eta$ and $|\tilde{\delta} - \delta| \leq 2\eta$.

The following lemma, whose proof is in OA.3.2, shows that there is an interval of

γ such that J and K can be tailored to obtain shares μ^{D_1} , μ^P , and μ^S within η of δ , $\bar{\mu}^P$, and $\bar{\mu}^S$, respectively, when noise is sufficiently small. (Consequently, the share μ^D must be within 2η of δ .) Moreover, the γ interval can be taken so that the incentives of supercooperators are satisfied.

Lemma 13. *There are $0 < \underline{\gamma} < \bar{\gamma} < 1$ and $\bar{\varepsilon} > 0$ such that, for all $\gamma \in [\underline{\gamma}, \bar{\gamma}]$ and $\varepsilon_C, \varepsilon_D < \bar{\varepsilon}$, there is a $D_J P_K S_1 D_\infty$ strategy with a steady state whose shares satisfy $|\mu_1^D - \delta|, |\mu^P - \bar{\mu}^P|, |\mu^S - \bar{\mu}^S| \leq \eta$, and are such that the $(C|D)_{J+K}$ constraint in Lemma 12 is satisfied.*

The left-hand side of the $(C|C)_J$ constraint in Lemma 12 converges uniformly to $\mu^S/(1 - \mu^{D_1})l$ as $\varepsilon \rightarrow 0$ for all $\gamma \in [\underline{\gamma}, \bar{\gamma}]$, $|\mu^{D_1} - \delta|, |\mu^P - \bar{\mu}^P|, |\mu^S - \bar{\mu}^S| \leq \eta$. Because $(\bar{\mu}^S - \eta)/(1 - \delta + \eta)l > g$, this means that $\bar{\varepsilon}$ can be chosen to be sufficiently small such that all these steady-state shares satisfy the $(C|C)_J$ constraint in Lemma 12 for all $\gamma \in [\underline{\gamma}, \bar{\gamma}]$ and $\varepsilon_C, \varepsilon_D < \bar{\varepsilon}$. This is similarly true for the $(D|D)_{J+K-1}$ constraint in Lemma 12, because the left-hand side of the corresponding inequality converges uniformly to $\gamma\mu^D/(1 - \gamma + \gamma\mu^D)l < l$ as $\varepsilon \rightarrow 0$ for all $\gamma \in [\underline{\gamma}, \bar{\gamma}]$, $|\mu^{D_1} - \delta|, |\mu^P - \bar{\mu}^P|, |\mu^S - \bar{\mu}^S| \leq \eta$.

Thus there are $0 < \underline{\gamma} < \bar{\gamma} < 1$ and $\bar{\varepsilon} > 0$ such that equilibria with shares μ^P, μ^S satisfying $|\mu^P - \bar{\mu}^P|, |\mu^S - \bar{\mu}^S| \leq \eta$ (and thus $\mu^C \geq 1 - 2\delta$) exist whenever $\gamma \in [\underline{\gamma}, \bar{\gamma}]$ and $\varepsilon_C, \varepsilon_D < \bar{\varepsilon}$.

A.3.3.3 Proof of Coordination-Proofness

We show that the grim trigger equilibria analyzed in A.3.3.1 and the $D_J P_K S_1 D_\infty$ equilibria analyzed in A.3.3.2 are coordination-proof. In any such equilibrium, (C, C) is played in every match where neither player has a defector record. By a similar argument to the proof of Lemma 2, the play in these matches is coordination-proof. Thus, we need only consider play in matches with a defector. Note that in equilibria generated by either grim trigger or $D_J P_K S_1 D_\infty$ strategies, the expected continuation value of a defector is weakly higher from playing D than from playing C . Since D is strictly dominant in the stage game, it follows that D is strictly dominant in the

augmented game for any defector. Thus, the prescribed action profile (D, D) in a match involving a reciprocator and a defector is the only equilibrium in the corresponding augmented game. Likewise, the prescribed action profile (C, D) in a match involving a supercooperator and a defector is the only equilibrium in the corresponding augmented game. We conclude that play in all matches is coordination-proof.