

Steady-State Equilibria in Anonymous Repeated Games, I: Trigger Strategies in General Stage Games*

Daniel Clark, Drew Fudenberg, and Alexander Wolitzky

August 30, 2019

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. We then focus on “trigger-strategy” equilibria. When the updating of a player’s record can depend on the actions of both players in a match, steady-state equilibria in trigger strategies can support the play of a wide range of actions, including any action that Pareto-dominates a static Nash equilibrium. When updates can depend only on a player’s own actions, fewer actions can be supported by steady-state equilibria. We provide sufficient conditions for trigger equilibria to support a given action, along with somewhat more permissive necessary conditions. When players have access to a form of decentralized public randomization, the sufficient conditions expand to match the necessary conditions.

*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 the past play of their partners, and little to no information about the play of people with whom their partners previously interacted. Yet, in some situations groups nevertheless 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 anonymous random matching. Our goals are to better understand what sorts of information suffice for good community outcomes, and to provide a foundation for the analysis of relatively simple strategies that we think are descriptively plausible. To this end, we suppose there is a continuum of players, which rules out both the contagion equilibria of Kandori (1992) and Ellison (1994) and the block belief-free equilibria constructed in Deb, Sugaya, and Wolitzky (2018). We assume that each player has a geometrically distributed lifetime, that the total mass of players is constant (with an inflow of new players replacing those who exit), and that the time horizon is doubly infinite, so there is no commonly known start date or notion of calendar time on which the players can coordinate their play. We then investigate how steady-state cooperation can be supported by simple strategies under various sorts of “record systems,” which provide each player with some summary of the reputation or standing of their current opponent.

To place our work in context, recall that in the standard repeated game model, a fixed finite set of players interact repeatedly with a commonly known start date and a common notion of calendar time. When each player’s signals are sufficient to statistically identify the vector of their opponents’ actions, equilibria that support

¹Examples of such “community enforcement” or “indirect reciprocity” include Milgrom, North, and Weingast (1990), Greif (1993), Greif, Milgrom, and Weingast (1994), Bernstein (1992), and Clay (1997) on merchant coalitions; Ostrom (1990) and Ellickson (1991) on managing public resources; Klein and Leffler (1981), Resnick and Zeckhauser (2002), and Dellarocas (2005) on seller reputation; Klein (1992), Padilla and Pagano (2000), and Bhaskar and Thomas (2018) on credit ratings; Friedman and Resnick (2001) on online ratings; and Fearon and Laitin (1996) on ethnic conflict.

cooperation usually exist when players are patient, but the most efficient equilibria are typically “complicated” if there is any noise in the monitoring structure. This model seems natural for studying some long-term relationships with well-defined start dates among a relatively small number of relatively sophisticated players, such as business partnerships or collusive agreements among firms. However, repeated games have also been used to model cooperation in large populations, as in the references in footnote 1. For these applications the assumptions of a fixed population, a common start date, and common calendar time seem less appropriate.²

In our model, when a pair of players match, each of them observes the other’s *record* before taking an action. A *record system* then updates the players’ records based on their current records and the actions they choose. We allow these systems to be stochastic, where the noise is due to either recording errors or to differences between a player’s intended action and the action that is implemented.

To begin, we prove that steady states exist for fairly general record systems, which allow the update of a player’s record to depend on both their own action and record as well as on the action and record of their opponent. The key requirement for this existence theorem is that the record update function be “finite partitional,” which means that for a given player record, there is a finite partition of the record space such that the update function depends on the opponent’s record in the same way for all opponent records in a given partition element. Along with our restriction to finite stage games, this assumption implies that for any given record, the record update function has a finite domain. We use this property together with geometric player lifespans to establish existence.

We then examine what sorts of actions can be *fully supported* by various record systems, meaning that there are equilibria in which the action is played by almost everyone when lifetimes are long and there is little noise. Here we restrict attention

²Even in repeated games with fixed partners, laboratory studies suggest that many subjects use fairly simple strategies. See e.g. Fudenberg, Rand, and Dreber (2012) and the survey by Dal Bó and Fréchet (2018).

to equilibria of a simple and plausible form, which we call *trigger systems*. These are pairs of strategy and record system such that the records can be partitioned into two classes—“good” and “bad”—where the prescribed play in a match depends only on the class of each player’s record, and the bad class is absorbing. This restriction lets us handle the set of possible equilibrium outcomes for general stage games; our companion paper Clark, Fudenberg, and Wolitzky (2019) analyzes more general strategies in the prisoner’s dilemma. We also show that steady-state equilibria in trigger strategies have desirable stability and convergence properties that cannot be guaranteed for general strategies, which further motivates the restriction to trigger strategies in this paper.

The record systems we start with are *second-order records*, which allow record updates to depend on the player’s own record and action as well as the partner’s action, but not the partner’s record. Thus, a second-order record system is able to update a player’s record based solely on the outcomes of their own interactions, without needing to know about the outcomes of their partner’s past interactions. In this sense, second-order records are more informationally robust than more general interdependent records. With second-order records, so long as players in good standing take different actions than players in bad standing when matched with good-standing opponents, trigger strategies can penalize players who deviate against good-standing partners by switching them to bad standing, without penalizing good-standing players for punishing bad-standing opponents. For this reason, second-order records enable trigger strategies to fully support a wide range of actions, including any action that Pareto-dominates a static Nash equilibrium. To show this, we use strategies of the following form: Players in good standing play the target action a with each other, and play a punishment action b against bad-standing players. Bad-standing players play a static equilibrium with each other and a best response c to the punishment action when facing good-standing opponents. Players start out in good standing and stay that way until they have accumulated sufficiently many “black marks,” where black marks are earned each time their outcome is recorded as anything other than a versus a or b versus c . The proof shows how to choose the threshold number K of black marks as a function of the

expected lifetime and the noise level to create a sequence of equilibria that converges to everyone playing a as the players' expected lifetimes grow long and the noise vanishes.³

We then turn to *first-order record systems*, which require that a player's record is updated based only on their own record and action. These record systems cannot support as many actions, and in particular cannot support a Nash-threat folk theorem, because first-order records cannot distinguish “justified” deviations from the target profile (a, a) from “unjustified” ones. For example, in the prisoner's dilemma, if a player is penalized for playing *Defect* against *Cooperate*, they must also be penalized for playing *Defect* against *Defect*. This makes it much harder to support a steady state with a high share of players in good standing.

We provide a sufficient condition for an action to be fully supported with first-order records using trigger strategies that require the punishment action b to be an “unprofitable deviation” against the target action a (that is, $u(a, a) > u(b, a)$) and moreover to be a better response than a against an action c that is a best response to b . Of course, such an action b is not guaranteed to exist; the sufficient condition for our proof is that one does.⁴ In our construction, players acquire black marks any time they are recorded as playing anything other than a or b (regardless of their partners' play), and switch permanently to bad status once they have a sufficient number of black marks. Once again, the heart of the proof is showing that there is a way to tune the threshold level of black marks to the parameters of the system so that there is a limit of equilibria in which everyone plays a .

We then provide a less restrictive sufficient condition for actions to be fully supported when records and strategies can depend on “personal public randomizations,” a decentralized form of public randomization. Finally we show that these less restrictive sufficient conditions are in fact necessary for the action to be fully sup-

³A subtlety is that the target action a must be distinct from the bad-standing best response action c , or else good-standing players may undetectably deviate from b to a against bad-standing opponents. This explains why trigger strategies cannot always support every individually rational action.

⁴Special cases of unprofitable punishments include “avoiding” the opponent or “burning money” to render a potentially tempting action unprofitable. We will discuss these special cases in some detail.

ported (barring knife-edge cases). Moreover, when these randomizations are available the strategies no longer need to keep a running tally of a player’s black marks; it is enough to know whether a player is in good or bad standing.

1.1 Related Work

In finite population models, Rosenthal (1979) and Rosenthal and Landau (1979) introduced the study of repeated games with random matching. Rosenthal (1979) considered the special case of first-order information where players observe only their current opponent’s most recent action, and established the existence of Markovian equilibria (which we call pairwise-public equilibria).⁵

Kandori (1992) and Ellison (1994) show that cooperation in the prisoner’s dilemma can be enforced by contagion equilibria when players have no information at all about each other’s past actions. These equilibria cannot exist in continuum-population models, or when a finite population is large compared to the discount factor. Nowak and Sigmund (1998) and many subsequent papers study the enforcement of cooperation using image scoring, which means that each player has first-order information about their partner, but conditions their action only on their partner’s record and not on their own record. These strategies are never a strict equilibrium, and are typically unstable in environments with noise (Panchanathan and Boyd, 2003).

In models with a continuum population, Okuno-Fujiwara and Postlewaite (1995) show that records that track a player’s “status” based on the actions and status levels of their opponents permit a folk theorem in the absence of noise.⁶ Takahashi (2010) shows how cooperation can be supported in the prisoner’s dilemma when players observe their partner’s entire past history of actions—all first-order information—but no higher-order

⁵Rosenthal and Landau (1979) study two particular record systems in an asymmetric battle of the sexes game. Their “comparative records” update based on both players’ records; their second, simpler, model has first-order records.

⁶Sugden (1986) and Kandori (1992) prove related results for finite populations. Steady-state equilibria in models with “status” also appear in the literature on fiat money following Kiyotaki and Wright (1993) and Kocherlakota (1998).

information. This construction relies on belief-free mixed strategies, which we will rule out. Takahashi also shows how cooperation can be supported in strict equilibrium when the prisoner’s dilemma stage game is strictly supermodular: that is, the loss from cooperation is lower when the partner cooperates. Heller and Mohlin (2018) suppose players observe a finite sample of their current partner’s past actions, which is a particular form of first-order information. Players live infinitely long and are completely patient, and are restricted to stationary strategies that depend only on their samples of their partner’s actions and not on their own histories. Their paper assumes a small fraction of players are commitment types, so that the partner’s past actions are a noisy signal of their type and thus of how they are likely to play today.⁷

Our interest in “simple” strategies in repeated games is shared with papers that study various notions of simplicity in fixed-partner repeated games, e.g. Rubinstein (1986), Abreu and Rubinstein (1988), Möbius (2001), Joe et al. (2012), and Compte and Postlewaite (2015). Repeated games with overlapping generations of players have also been studied by e.g. Cremer (1986), Kandori (1992), Salant (1991), and Smith (1992); however, these papers are less directly relevant as they consider non-anonymous players. Finally, the random matching model of Fudenberg and He (2018, 2019) is similar to ours in several respects, including countably infinite agent histories, geometric agent lifetimes, and a doubly infinite time horizon.

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. The time horizon is doubly infinite. When two players match, they play a finite symmetric game with action space A and payoff function $u : A \times A \rightarrow \mathbb{R}$.

⁷Bhaskar and Thomas (2018) study first-order information in a sequential-move game with one-sided incentive constraints.

This section presents our model of records, states, and steady-state equilibria with general pairwise-public strategies; in later sections of the paper we specialize to equilibria in trigger strategies.

2.1 Record Systems

Every player carries a *record*, and when two players meet, each sees 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 triple (R, \mathcal{D}, ρ) comprised of a countable set R (the *record set*), a countable set $\mathcal{D} \subset R^2 \times A^2$ (the *update domain*), and a function $\rho : \mathcal{D} \rightarrow \Delta(R)$ (the *update rule*).

We assume that 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.

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 refer to the general case where $\mathcal{D} = R^2 \times A^2$ 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 two prior examples of finite-partitional interdependent record systems.

Many simple and realistic record systems fall into a more restricted class, where a player’s update does not depend directly on their opponent’s record, so that the player’s record can be updated even when the opponent’s record is not available.

We say a record system is *second-order* if $\mathcal{D} = R \times A^2$, so the update can depend on both players’ actions but only on the player’s own record. With a second-order

system, a player’s record can be computed based only on their own history of stage-game outcomes.⁸ A record system is *first-order* if the update can depend only on the player’s own action and record as in e.g. Nowak and Sigmund (1998), Takahashi (2010), and Heller and Mohlin (2018). Here $\mathcal{D} = R \times A$. We consider second-order systems in Section 4 and first-order systems in Section 5.

Note that we allow stochastic update rules. This can represent errors in recording, and it can also model imperfect implementation of the intended action. To apply the latter interpretation with first-order records, 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 seen. However, since there is a continuum of players, only the player’s current record and that of their current partner matter for the player’s current payoff, and only the player’s own record will matter in the future. For this reason, all 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 pure pairwise-public strategies as functions $\mathbf{s} : R \times R \rightarrow A$, with the convention that the first coordinate is the player’s own record and the second coordinate is that of the partner, and similarly write mixed pairwise-public strategies as functions $\boldsymbol{\sigma} : R \times R \rightarrow \Delta(A)$. For simplicity, we also assume that all players use the same strategy. Note that every strict equilibrium in a symmetric, continuum-

⁸If the record system tracks the time periods in which each pair of actions occurs, and contrary to our assumptions it is “centralized” in the sense of having access to all players’ records, then second-order records could be used to compute and track the “status levels” implied by a more general interdependent record system. However, this is not true for the decentralized records we consider (where the update domain is a subset of $R^2 \times A^2$), and even with centralized records a simple count of the number of times each action profile occurred would not suffice to reproduce interdependent records. Kocherlakota (1998) makes a similar point in the context of macroeconomic matching models.

population model is symmetric, so this restriction would be without loss if we restricted attention to strict equilibria. For generic stage games, the equilibria we construct in this paper are strict if the stage game admits a strict and symmetric Nash equilibrium.

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

Given a record system \mathcal{R} and a pairwise-public strategy σ , we can define an update map as follows: First, when all players use strategy σ , denote the distribution over next-period records of a player with record r who meets a player with record r' by

$$\phi_{r,r'}(\sigma) = \sum_a \sum_{a'} \sigma(r, r')[a] \sigma(r', r)[a'] \rho(r, r', a, a') \in \Delta(R).^9$$

Then, the *update map* $f_\sigma : \Delta(R) \rightarrow \Delta(R)$ is given by

$$\begin{aligned} f_\sigma(\mu)[0] &:= 1 - \gamma + \gamma \sum_{r'} \sum_{r''} \mu_{r'} \mu_{r''} \phi_{r',r''}(\sigma)[0], \\ f_\sigma(\mu)[r] &:= \gamma \sum_{r'} \sum_{r''} \mu_{r'} \mu_{r''} \phi_{r',r''}(\sigma)[r] \text{ for } r \neq 0. \end{aligned}$$

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

Theorem 1.

- *Under any first or second-order record system and any pairwise-public strategy, a steady state exists.*
- *Under any finite-partitional interdependent record system and any pairwise-public*

⁹This equation applies for general interdependent records. To apply it to the other sorts of records we embed their domains in the general one in the obvious way: for second-order records, the dependence of ρ on its second argument is trivial; for first-order records, the dependence on the fourth argument is also trivial.

strategy, a steady state exists.

- *For interdependent record systems that are not finite-partitional, a steady state may fail to exist.*

The proof is in A.1 of the Appendix; all other omitted proofs can be found in either the Appendix (A) or the Online Appendix (OA).¹⁰ In outline, we relabel records so that two players with different ages can never share the same record, define a set \bar{M} that contains all feasible distributions over records, and show that \bar{M} and the update map jointly satisfy the conditions for a fixed-point theorem. Intuitively, the combination of the finite domain of the record-update function (due to finiteness of the stage game and, for interdependent record systems, the finite-partition property) and geometrically distributed lifetimes imply that most players have records in a finite subset of the set $\bigcup_{t \leq T} R(t)$ for bounded T , so \bar{M} resembles a finite-dimensional space, and in particular 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. When instead the record-update function does not have finite domain, the update map can map any state to one with more weight in the upper tail in such a way that no steady state exists. The proof shows that this is the case for example if 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.¹¹ Intuitively, this multiplicity corresponds to different initial conditions at time $t = -\infty$.

¹⁰Fudenberg and He (2018) prove that steady states exist in their random matching model with geometric lifetimes and countably many records; in their model agents do not observe each other's records, but their own record is updated based on both players' actions.

¹¹For 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 so for the remainder of their lifetime. Depending on the parameters γ and ε , there can be between one and three steady states.

2.3 Steady-State Equilibria

We focus on those 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

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

Next, denote the probability that a player with record r today has record r' t periods from now (assuming the player is still alive then) by $\phi_r^t(\sigma, \mu)[r']$. This is defined recursively by

$$\phi_r^1(\sigma, \mu)[r'] = \sum_{r''} \mu_{r''} \phi_{r, r''}(\sigma)[r']$$

and, for $t > 1$,

$$\phi_r^t(\sigma, \mu)[r'] = \sum_{r''} (\phi_r^{t-1}(\sigma, \mu)[r'']) (\phi_{r''}^1(\sigma, \mu)[r']).$$

The continuation value of a player with record r is then given by

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

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

A pair (σ, μ) is an *equilibrium* if μ is a steady-state under σ and, for each own record r and opponent's record r' , we have

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

for all a such that $\sigma(r, r')[a] > 0$. Thus, a player's objective is to maximize their expected undiscounted lifetime payoff. An equilibrium is *strict* if the argmax is unique for all pairs of records (r, r') , so that the player has a strict preference for following the

equilibrium strategy. As noted above, every strict equilibrium is pairwise-public, pure, and symmetric.

Corollary 1. *Under any first or second-order record system, an equilibrium exists. The same is true for any finite-partitional interdependent record system.*

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

In contrast, the existence of a strict equilibrium is not guaranteed. A sufficient condition for strict equilibrium existence is for the stage game to have a strict and symmetric Nash equilibrium.

Corollary 2. *Under any first or second-order record system, a strict equilibrium exists if the stage game has a strict and symmetric Nash equilibrium. The same is true for any finite-partitional interdependent record system.*

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

3 Trigger Strategies and Noisy Counts

3.1 Trigger Strategies

Given a steady state μ for pure pairwise-public strategy \mathbf{s} and record system \mathcal{R} , record r' is *reachable from* r if there is a t with $\phi_{s,\mu}^t(r)[r'] > 0$.

Definition 2. *The record system (\mathcal{D}, R, ρ) and pure strategy \mathbf{s} are a **trigger system** if*

1. *there are maps $\psi : R \rightarrow \{G, B\}$ and $\xi : \{G, B\} \rightarrow A$ such that $\mathbf{s}(r, r') = \xi(\psi(r), \psi(r'))$, and*

2. if $\psi(r) = B$ then $\psi(r') = B$ for all r' that are reachable from r .

When this is the case we say that \mathbf{s} is a **trigger strategy**.

Condition (1) says that the records can be partitioned into two classes so that the prescribed play in a match depends only on the class of each player's record, and condition (2) says that the "bad" class B is absorbing. We call these "trigger strategies" because of their resemblance to the trigger strategies used by e.g. Friedman (1971), but those strategies only prescribed two distinct actions, one on the path of play and one as a response to deviations. Here a trigger strategy can prescribe up to four distinct actions, one for each of the four ordered pairs of classes. The actions for each pair are not required to be distinct; for example, a game with a pure strategy Nash equilibrium has an equilibrium in trigger strategies for any record system, as we can specify that all records are in B and all records play the (same) Nash equilibrium action.

The restriction that records can be partitioned into two classes, with play depending on records only via their class, does involve a loss of generality, even in games with only two actions. For example, in Clark, Fudenberg, and Wolitzky (2019), for some parameters of the prisoner's dilemma the only equilibrium with two classes is *Always Defect*, while equilibria with three classes can yield more efficient outcomes. The requirement that one state be absorbing is also restrictive, as in some cases "cyclic" strategies with two classes can outperform trigger strategies. Nevertheless, the space of trigger strategies is rich enough to allow efficient equilibria in some cases. Unlike the space of all record-dependent strategies, it is also sufficiently tractable that we will be able to give necessary and sufficient conditions for these strategies to support steady-state equilibria where almost everyone plays a given action.

Note that, since the bad class is absorbing, when two players with bad records meet they must play a static Nash equilibrium. To respect our restriction to symmetric, pure strategies, our positive results thus assume that the stage game admits a symmetric, pure equilibrium; this can be relaxed by letting bad-standing players play a mixed equilibrium.

3.2 Noisy Counts

Our sufficient conditions use trigger strategies that operate on a simple class of record systems, where the records simply count the number of times each action profile (second-order system) or own action (first-order system) occurred in a player’s history. Let $n = |A|$, order the action profiles in A^2 as $((a, a')_1, \dots, (a, a')_{n^2})$, and let $e_{(a, a')}$ be the unit vector corresponding to (a, a') ; similarly, order the actions in A as (a_1, \dots, a_n) , and let e_a be the unit vector corresponding to a .

Definition 3. *A second-order record system is a **noisy count** if the record r can be written as a vector $(r_1, r_2, \dots, r_{n^2})$ where new players have record $(0, 0, \dots, 0)$ and $\rho(r, a, a') = (1 - (n^2 - 1)\varepsilon)(r + e_{(a, a')}) + \sum_{(\bar{a}, \bar{a}') \neq (a, a')} \varepsilon(r + e_{(\bar{a}, \bar{a}')})$, for some $\varepsilon \in (0, 1)$.*

*A first-order record system is a **noisy count** if the record r can be written as a vector (r_1, r_2, \dots, r_n) where new players have record $(0, 0, \dots, 0)$ and $\rho(r, a) = (1 - (n - 1)\varepsilon)(r + e_a) + \sum_{\bar{a} \neq a} \varepsilon(r + e_{\bar{a}})$, for some $\varepsilon \in (0, 1)$.*

With noisy second-order counts, the intended action profile can be mis-recorded as some other profile. With noisy first-order counts, the intended action can be mis-recorded as another action. For simplicity, we state the noisy count definitions with exactly the same level of noise on each action, but we have verified that the constructions used in our positive results (Theorems 2, 3, and 4) are all valid for an open interval of noise levels.

Noisy counts are very special record systems. However, since partners can always agree not to condition their behavior on jointly observed random variables, all of our results for noisy counts hold a fortiori for any systems where records convey additional information, for example information about the time path of the partner’s actions.

The trigger strategies we will construct on noisy count record systems all lie in the following simple class, where a record is classified as “good” if and only if the sum of the counts of some actions or action profiles is less than a fixed cutoff score $K \in \mathbb{N}$.

Definition 4. *Given a noisy count record system, a trigger system is a **count trigger system** if there exist indices m_1, \dots, m_M and a cutoff score $K \in \mathbb{N}$ such that, for*

each record $r = (r_1, \dots, r_{n^2})$ (for second-order records) or (r_1, \dots, r_n) (for first-order records), we have $\phi(r) = G$ if and only if $\sum_{m=m_1}^{m_M} r_m < K$. Strategies in a count trigger system are **count trigger strategies**.

The set of equilibria depends on both the amount of noise in the system and the players' expected lifetimes. In what follows, we ask which actions can be played by almost the entire population when there is vanishingly little noise and players live a long time. To make this precise, for any $a \in A$, let $\bar{\mu}_{\gamma, \varepsilon}^a$ be the supremum of the share of players playing a in a steady-state strict equilibrium for parameters (γ, ε) .

Definition 5. Action a can be **fully supported in the limit** if $\lim_{(\gamma, \varepsilon) \rightarrow (0, 0)} \bar{\mu}^a(\gamma, \varepsilon) = 1$.

4 Second-Order Records

This section shows how to fully support a wide range of actions with noisy second-order counts. 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.

Theorem 2. Fix an action a and suppose there exist actions b and c such that $c \neq a$, c is a best-response to b , and $u(a, a) > u(c, b)$. Then action a can be fully supported in the limit by trigger strategies with noisy second-order counts.

The requirements that $a \neq c$, c is a best-response to b , and $u(a, a) > u(c, b)$ together imply that $a \neq b$; however, the theorem allows $b = c$. For example, in the prisoner's dilemma, trigger strategies in which $a = \textit{Cooperate}$ and $b = c = \textit{Defect}$, so that the intended action profile is $(\textit{Cooperate}, \textit{Cooperate})$ when two good-standing players match but is otherwise $(\textit{Defect}, \textit{Defect})$, can fully support cooperation in the limit.

More generally, if the stage game admits a pure Nash equilibrium (d, d) , taking $b = c = d$ implies that any action a such that $u(a, a) > u(d, d)$ can be fully supported in the limit; this gives a “Nash-threat folk theorem.” In addition, say that an action a is “not uniquely optimal against minmax” if there exists a pure strategy minmax action b against which a is not the unique best response. Theorem 2 implies that any strictly individually rational action that is not uniquely optimal against minmax can be fully supported in the limit.

Here is an outline of the proof of Theorem 2, which is in A.2.2. Let (d, d) denote a pure static Nash equilibrium. The proof defines a *score* $k \in \mathbb{N}$ for each player, which equals the number of times their outcome was recorded as anything other than (a, a) or (b, c) (i.e., the player playing b against an opponent who plays c). Players are in good standing for k strictly less than some cutoff score K , and are otherwise in bad standing. Players in good standing play a when they are matched with a fellow good-standing player, and play b against bad-standing players. Players in bad standing play c when matched with good-standing players, and play d against fellow bad-standing players.

To prove the theorem, it suffices to find a cutoff score K as a function of γ and ε such that the corresponding strategy profile is an equilibrium and the steady-state share of good-standing players is close to 1 when $(\gamma, \varepsilon) \approx (1, 0)$. Note that K cannot be fixed independent of (γ, ε) ; for example, for any fixed K , if $1 - \gamma$ is much smaller than ε then the steady-state share of good-standing players will be close to 0. However, K also cannot be too large, as otherwise newborn players (who are the farthest from reaching bad standing) will deviate.

To show that an appropriate cutoff exists, define a function $\alpha : (0, 1)^2 \rightarrow (0, 1)$ by

$$\alpha(\gamma, \varepsilon) = \frac{\gamma\nu(\varepsilon)}{1 - \gamma + \gamma\nu(\varepsilon)}, \quad (1)$$

where $\nu(\varepsilon) = (n^2 - 2)\varepsilon$ is the probability that a good-standing player’s count increases in a given period, so $\alpha(\gamma, \varepsilon)$ gives the probability that a good-standing player lives to see their score increase by at least 1.

Lemma 1. *There is a constant $z > 0$ (independent of γ and ε) such that the trigger strategy with tolerance K has an equilibrium with share of good-standing players μ^G if the feasibility constraint*

$$\mu^G = 1 - \alpha(\gamma, \varepsilon)^K$$

and the incentive constraint

$$\frac{1 - \mu^G}{\nu(\varepsilon)} (\mu^G (u(a, a) - u(c, b)) + (1 - \mu^G) (u(b, c) - u(d, d))) > z$$

are satisfied.

The proof of Lemma 1 is in A.2.1. The feasibility constraint comes from calculating the steady-state shares μ_k for the trigger strategy and then setting $\mu^G = \sum_{k=0}^{K-1} \mu_k$. Intuitively, because $\alpha(\gamma, \varepsilon)^K$ is the probability that a newborn player lives to see their score exceed K (at which point they enter bad standing), it also equals $1 - \mu^G$, the share of bad-standing players in steady-state. The incentive constraint comes from solving the value functions V_k and then evaluating the incentives of the good-standing players to follow the prescribed play. For a good-standing player with score 0, the expected increase in the number of periods they will spend in bad standing rather than good standing when their score increases by 1 equals the ratio of $1 - \mu^G$, the probability a player with score 0 ultimately reaches bad standing, and $\nu(\varepsilon)$, the probability a good-standing player's score increases in a given period. Since $\mu^G (u(a, a) - u(c, b)) + (1 - \mu^G) (u(b, c) - u(d, d))$ is the difference in good-standing and bad-standing flow payoffs, when the product of this term and $(1 - \mu^G)/\nu(\varepsilon)$ is sufficiently high, good-standing players with score 0 (who are the players most tempted to cheat) play according to the prescribed strategy both when matched against good-standing players and when matched with bad-standing players.

The proof of Theorem 2 then proceeds by setting the tolerance level $K(\gamma, \varepsilon)$ to be the smallest integer larger than $\ln(\eta)/\ln(\alpha(\gamma, \varepsilon))$ for fixed $\eta \in (0, 1)$. Note that this is the smallest tolerance level such that the corresponding μ^G is larger than $1 - \eta$. The

proof shows that this $K(\gamma, \varepsilon)$ is small enough that the incentive constraints of score 0 good-standing players are satisfied. In particular, as $(\gamma, \varepsilon) \rightarrow (1, 0)$, the ratio $(1 - \mu^G(\gamma, \varepsilon))/\nu(\varepsilon)$ blows up, which leads to the incentive constraints being satisfied in the limit when η is sufficiently small.

Okuno-Fujiwara and Postlewaite (1995)'s Theorem 1 shows that with “status” (a form of interdependent record) any actions that Pareto-dominate the pure-strategy minmax payoffs can be implemented in equilibrium without noise. Their proof uses grim trigger strategies and hence is not robust to noise. Nonetheless, Theorem 2 shows that the theorem’s conclusion does not require interdependent records (provided $a \neq c$), and extends to cases with overlapping generations and noise. This establishes that a range of equilibrium behavior can arise in large populations under more general conditions than past work had indicated.

To illustrate the need for $a \neq c$ in Theorem 2, consider the following game.

	A	B
A	2, 2	1, 3
B	3, 1	0, 0

Here action A cannot be supported by trigger strategies, as bad-standing players must play their best response A to the punishing action B , but then good-standing players will play A against bad-standing opponents as well as good-standing ones. Using similar arguments as in the proof of Theorem 2, we can show that A can however be fully supported in the limit with “cyclic” strategies, where the promise of eventually returning to good standing motivates bad-standing players to play B ; indeed, for any stage game, any action that Pareto-dominates the pure strategy minmax payoffs can be fully supported in the limit by cyclic strategies with second-order records.¹² In addition, essentially the same argument as the proof of Theorem 2 shows that any action that Pareto-dominates pure strategy minmax can be fully supported in the limit by trigger strategies with interdependent records: in this case, the $a \neq c$ requirement can be dropped, because the record system can distinguish between (equilibrium) plays

¹²The proof is available from the authors upon request.

of (a, a) against a good-standing opponent and (deviant) plays of (a, a) against an opponent in bad standing.

5 First-Order Records

Now we turn to first-order record systems, where the updating of a player’s record depends only on their own play. As discussed in the introduction, such records cannot support as many actions, and in particular the folk theorem fails.¹³ The space of all possible strategies using any conceivable first-order record system is very large, and completely characterizing when a given action can be fully supported using general strategies is beyond the scope of this paper. We instead restrict attention to trigger strategies, a simple and tractable class. We provide two sets of sufficient conditions for trigger strategies to fully support a given action: one with the noisy first-order counts defined above, and a second when these records are augmented with “personal public randomization” (PPR). We then show that the sufficient conditions with PPR are also necessary for an action to be fully supported with trigger strategies.

5.1 Unprofitable Punishments

An important case where trigger strategies can fully support action a in the limit is when there exists a punishing action b and a 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 the punishing action is a better response to c than the target action). We say that in this case b is an “unprofitable punishment” for a .

¹³The failure of the folk theorem with first-order records does not depend on restricting to trigger strategies. Clark, Fudenberg, and Wolitzky (2019) show that the only pure-strategy equilibrium in the prisoner’s dilemma is *Always Defect* when the gain from playing *Defect* is strictly greater when the opponent *Cooperates* rather than *Defects*. In a related model, Takahashi (2010) shows that cooperation can be fully supported in the prisoner’s dilemma if belief-free mixed equilibria are allowed. We exclude these equilibria since they may not be purifiable (Bhaskar, Mailath, and Morris, 2008) or evolutionarily stable (Heller, 2017).

Definition 6. Action b is an *unprofitable punishment* for action a if there exists a best response c to b such that

1. $u(a, a) > u(c, b)$,
2. $u(a, a) > u(b, a)$, and
3. $u(b, c) > u(a, c)$.

Theorem 3. *With noisy first-order counts, any action a for which there is an unprofitable punishment can be fully supported in the limit with count trigger strategies.*

The proof, which is in A.3.2, is similar to the proof of Theorem 2. Intuitively, since the punishment action b is not a profitable deviation from the target profile (a, a) , the fact that first-order records cannot distinguish justified plays of b from deviations does not undermine the players' incentives. In particular, a player's score does not need to increase when they play b .

To sketch the proof in more detail, let (d, d) be a static pure equilibrium. Let a player's score k be the number of times their action was recorded as anything other than a or b . Players are in good standing for $k < K$ for some cutoff score K , and are otherwise in bad standing. Players in good standing play a when they are matched with a fellow good-standing player, and play b against bad-standing players. Players in bad standing play c when matched with good-standing players, and play d against fellow bad-standing players.

Define the function $\beta : (0, 1)^2 \rightarrow (0, 1)$ by

$$\beta(\gamma, \varepsilon) = \frac{\gamma(n-2)\varepsilon}{1-\gamma+\gamma(n-2)\varepsilon}. \quad (2)$$

Note that $\beta(\gamma, \varepsilon)$ gives the probability that a good-standing player lives to see their score increase by at least 1. Let y represent the greatest one-shot gain that a good-standing player can obtain by playing their static best response to either a or c .¹⁴

¹⁴That is, $y = \max \{ \max_x u(x, a) - u(a, a), \max_x u(x, c) - u(b, c) \}$.

Lemma 2 adapts Lemma 1’s result about general punishments with second-order records to unprofitable punishments with first-order records (proof in A.3.1).

Lemma 2. *The trigger strategy with tolerance K has an equilibrium with share of good-standing players μ^G if the feasibility constraint*

$$\mu^G = 1 - \beta(\gamma, \varepsilon)^K$$

and the incentive constraint

$$\frac{1 - n\varepsilon}{(n - 2)\varepsilon}(1 - \mu^G)(\mu^G(u(a, a) - u(c, b)) + (1 - \mu^G)(u(b, c) - u(d, d))) > y$$

are satisfied.

The proof of Theorem 3 then proceeds by setting the tolerance level $K(\gamma, \varepsilon)$ to be the smallest integer larger than $\ln(\eta)/\ln(\beta(\gamma, \varepsilon))$. Similar to the proof of Theorem 2, the proof then shows that as $(\gamma, \varepsilon) \rightarrow (1, 0)$, this tolerance level induces a population share of good-standing players that is larger than $1 - \eta$, but is not so large that the good-standing players’ incentive constraints are violated.

We now present two leading cases in which an action a has an unprofitable punishment, and can therefore be fully supported in the limit with count trigger strategies.

The first case arises when $b = c = d$ in the definition of an unprofitable punishment, so that (b, b) is a stage-game Nash equilibrium, and trigger strategies support (a, a) through Nash reversion.

Corollary 3. *If there is an action b such that (b, b) is a Nash equilibrium of the stage game and $u(a, a) > \max\{u(b, a), u(b, b)\}$, then a can be fully supported in the limit with count trigger strategies.*

For example, suppose the stage game is a prisoner’s dilemma with an *exit option* E . In this game, when either player plays E , both players receive the same payoff, which is less than the cooperative payoff $u(C, C)$ but more than the “sucker’s payoff” $u(C, D)$,

and not more than the non-cooperative payoff $u(D, D)$. Note that both (E, E) and (D, D) are static equilibria, but E is not a profitable deviation against C , unlike D . Corollary 3 says that Nash reversion to (E, E) can fully support cooperation in the limit.

This example is closely related 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 has led Bowles and Gintis (2011), Boyd et al. (2003), and Gintis et al. (2003) (among others) to argue that the enforcement of human cooperation cannot be explained without appealing to social preferences. Other authors (Baumard (2010), Guala (2012)) argue that human cooperation is better explained by simply avoiding deviators, rather than actively punishing them. The fact that trigger strategies can support full limit cooperation in the prisoner’s dilemma with second-order records but not with first-order records supports the argument that the inability to distinguish justified and unjustified plays of *Defect* is a serious obstacle to cooperation in the prisoner’s dilemma. However, this obstacle evaporates when a simple exit option is added to the game, consistent with the position of Baumard and Guala

The second special case arises when “money-burning” is available, in that players can observably reduce their own utility by any amount, simultaneously with 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 in the sense of Definition 6. 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 money that can be burned along with playing b to make this punishment unprofitable.

5.2 Personal Public Randomizations

We now show that enriching the information structure to allow a simple class of randomizing devices, which we call *personal public randomizations (PPR)*, expands the set of actions that can be fully supported in the limit by trigger strategies with noisy first-order counts. In the next subsection, we will show that any action outside this expanded set cannot be fully supported by trigger strategy for any noisy first-order record system (except for knife-edge cases). Thus, when PPR are available, we essentially completely characterize the set of actions that can be fully supported by trigger strategies with all noisy first-order record systems that include action counts.

With PPR, whenever a player's action is recorded, this action is associated in the player's record with an indelible random variable ω that has the uniform distribution on $[0, 1]$. Here, we assume that players observe the entire history of their opponent's noisily recorded actions along with the corresponding PPR, so that the record of these PPR can be used by the player's future partners to condition their play. PPR are thus a "decentralized" form of public randomization (since it attaches to each player's record separately), and lead to some simpler and cleaner results, similarly to how public randomization is often a useful simplification in fixed-partner repeated games. Covering the two extreme cases of no correlating devices and PPR, as we do here, seems likely to give a good sense of the boundary of what could be achieved with more realistic, imperfect forms of correlation.

Theorem 4. *With noisy first-order counts and PPR, trigger strategies can fully support action a in the limit if either*

- (a, a) is a Nash equilibrium of the stage game, or
- there exist actions b and c such that c is a best response to b and
 1. $u(a, a) > u(c, b)$,
 2. $u(b, a) - u(a, a) < u(b, c) - u(a, c)$,

3. $u(b, c) > u(a, c)$,
4. $u(b, a) - u(a, a) < u(a, a) - u(c, b)$.

Note that if b is an unprofitable punishment for a then Conditions 1–4 of Theorem 4 are satisfied: Conditions 1 and 3 of Theorem 4 are the same as Conditions 1 and 3 in the definition of an unprofitable punishment, and Conditions 2 and 4 of Theorem 4 are satisfied because Condition 2 in the definition of an unprofitable punishment implies that $u(b, a) - u(a, a)$ is negative, while Conditions 1 and 3 imply that $u(b, c) - u(a, c)$ and $u(a, a) - u(c, b)$ are both positive. Theorem 4 thus extends Theorem 3 to the case where b is a “slightly profitable” deviation, in that $u(b, a) - u(a, a)$ is positive but less than $\min\{u(b, c) - u(a, c), u(a, a) - u(c, b)\}$.

Turning to the proof, the case where (a, a) is a Nash equilibrium is immediate. The proof for the second case (in A.4) extends the trigger strategy construction of Theorem 3 by constructing strategies that use PPR to govern the way that records move through two “phases” labelled G (“good standing”) and B (“bad standing”). Players start out in G , and they transition to B with probability τ_a when they are recorded playing a , with probability τ_b when they are recorded playing b , and with probability χ when they are recorded playing any other action.

The proof proceeds by showing that it is feasible to target any share of good-standing players $\mu^G \in (0, 1)$ in the $(\gamma, \varepsilon) \rightarrow (1, 0)$ limit by appropriately choosing the transition probabilities $\tau_a(\gamma, \varepsilon)$, $\tau_b(\gamma, \varepsilon)$, and $\chi(\gamma, \varepsilon)$. This comes from the fact that with any fixed transition probability, good-standing players are very likely to eventually transition to bad standing when γ is large, so the desired share μ^G can be implemented by making the transition probabilities smaller as γ gets closer to 1.

Moreover, the proof shows that the transition probabilities can also be chosen so that the incentive constraints are satisfied in the limit when μ^G is sufficiently close to 1. Conditions 1 and 4 guarantee that when $\mu^G = 1$, the difference between the flow payoff to good standing and that to bad standing is larger than the payoff gain a good-standing player would experience by playing b when they are supposed to play a .

Since $u(b, c) > u(a, c)$ and $u(b, a) - u(a, a) < u(b, c) - u(a, c)$, when μ^G is sufficiently close to 1, we can choose the transition probability after playing b to be sufficiently larger than the transition probability after playing a to prevent good-standing players from playing b when they are supposed to play a , but also not so much larger that the good-standing players prefer to play a when they should play b . Furthermore, we can choose the transition probability after playing any $x \neq a, b$ to be much larger than the transition probability after either a or b , intuitively because it is much less likely that a good-standing player is recorded as playing some $x \neq a, b$ than either a or b when ε is very small. This deters the good-standing players from playing any action other than a or b , regardless of the increase in stage-game payoff they could achieve.

5.3 Necessary Conditions

This section shows that weak versions of the sufficient conditions of Theorem 4 are also necessary for trigger strategies to fully support action a for any “noisy” first-order record system (a much richer class than noisy first-order counts).

Definition 7. A first-order record system is **noisy** if for each record r and $i, j \in \{1, \dots, n\}$ there exist $q_i(r) \in \Delta(R)$ and $\nu_{i,j}(r) \in (0, 1)$ such that $\sum_i \nu_{i,j}(r) = 1$ and $\rho(r, a_j) = \sum_{i \in \{1, \dots, n\}} \nu_{i,j}(r) q_i(r)$.

Here $q_i(r)$ represents the distribution over records after “a recording of a_i is fed into the record system” and $\nu_{i,j}(r)$ is the probability that this occurs when a record- r player plays a_j . This noise can come from either the recording system itself (“recording errors”) or in the map from intended to realized actions (“implementation errors”).

Theorem 5. For any noisy first-order record system, if for every $\eta > 0$ there exist some γ and ε for which trigger strategies support share at least $1 - \eta$ of the population playing a , then there exist actions b, c , and d (not necessarily distinct), such that c is a best response to b , (d, d) is a Nash equilibrium of the stage game, and

1. $u(a, a) \geq u(c, b)$,

2. $u(b, a) - u(a, a) \leq u(b, c) - u(a, c)$,
3. $u(b, c) \geq u(a, c)$, and
4. $u(b, a) - u(a, a) \leq u(a, a) - u(c, b)$.

We now explain the specific necessary conditions required by Theorem 5, For more than half the population to play a in a trigger strategy equilibrium, the strategy must dictate that either good-standing players play a against fellow good-standing players, or bad-standing players play a against fellow bad-standing players. If (a, a) is not a Nash equilibrium, the latter is impossible, so good-standing players must play a against fellow good-standing players.¹⁵

Consider the class of trigger strategies in which good-standing players play a against fellow good-standing players and action b against bad-standing players, and bad-standing players play c against good-standing players and d against bad-standing players. Because the bad class is absorbing, c must be a best response to b , and (d, d) is a Nash equilibrium. To establish the remaining four conditions (assuming (a, a) is not a Nash equilibrium), it will be helpful to first show that $a \neq b$.

Lemma 3. *If trigger strategies can support an arbitrarily large share of the population playing a , then $b \neq a$.*

Proof. If instead $b = a$, any player can achieve a strictly higher stage-game payoff than $u(a, a)$ when matched with good-standing players, namely $\max_x u(x, a)$. Thus, this class of trigger strategies cannot sustain equilibria for arbitrarily high μ^G , because the expected lifetime payoff of newborn players becomes arbitrarily close to $u(a, a)$ as μ^G converges to 1, while players can secure a payoff arbitrarily close to $\max_x u(x, a) > u(a, a)$ as μ^G converges to 1. ■

¹⁵Note that the necessary conditions of Theorem 5 are always satisfied when (a, a) is a Nash equilibrium, as can be seen by taking $b = c = d = a$.

A preliminary result is that when there are noisy first-order records and (a, a) is not a Nash equilibrium of the stage game, there must be a positive share of bad-standing players ($\mu^G < 1$) in any trigger strategy equilibrium.¹⁶

Lemma 4. *In any trigger strategy equilibrium with noisy first-order records, $\mu^G < 1$.*

Proof. Suppose that $\mu^G = 1$. Because records are noisy, a player at a good-standing record could play their best response to a , and be assured of remaining in good-standing record for the rest of their life with probability 1. Thus, it is strictly better for a good-standing player to play their best response to a when matched with a fellow good-standing player rather than follow the prescribed trigger strategy. ■

The remainder of the proof of Theorem 5 (omitted details of which are given in A.5) establishes Conditions 1-4 in turn. For some intuition, Condition 1 is the limit as the share of good-standing players approaches 1 of the requirement that the flow payoff of a good-standing player exceeds that of a bad-standing player, which is obviously necessary for trigger strategies to be an equilibrium when (a, a) is not a static Nash equilibrium. Condition 2 is an increasing differences condition; it is necessary for a good-standing player to prefer to play a rather than b when the opponent plays a , while preferring b when the opponent plays c .

Conditions 3 and 4 are more subtle. Condition 3 says that b is a better response to c than a , and Condition 4 says that the difference in flow payoff between good-standing and bad-standing players exceeds the short-term gain from playing b rather than a against a . The rest of this subsection explains these conditions in more detail; it is not needed to understand what follows.

For simplicity, suppose that newborn players are in good standing and have the highest continuation values in the population (as is the case for count trigger strategies). Note that the continuation payoff V_0 of a newborn player equals the average payoff in

¹⁶The fact that equilibria with noisy records must have $\mu^G < 1$ is what leads to the necessary conditions in Theorem 5. The alternative hypothesis, “For any (possibly non-noisy) first-order record system, if for every $\eta > 0$ there exist some γ and ε for which trigger strategies support share $1 - \eta < \mu^G < 1$ of the population playing a ,” would give the same conclusion.

the population, because the expected fraction of a player's lifetime spent at record r is equal to the fraction of the population with record r (and there is no discounting, so both V_0 and the population-average payoff are given by $\sum_r \mu_r \pi_r$). The average payoff in the population is $\mu^G \pi^G + (1 - \mu^G) \pi^B$, where $\pi^G = \mu^G u(a, a) + (1 - \mu^G) u(b, c)$ is the flow payoff to good-standing players and $\pi^B = \mu^G u(c, b) + (1 - \mu^G) u(d, d)$ is the flow payoff to bad-standing players.

Since a good-standing player plays a when matched with a fellow good-standing player and b when matched with a bad-standing player, $V_0 = (1 - \gamma) \pi^G + \gamma \mu^G V_0^a + \gamma (1 - \mu^G) V_0^b$, where V_0^a and V_0^b are the expected continuation payoffs of a newborn player after playing a and b , respectively. Since $V_0 \geq V_r$ for all $r \in R$, both $V_0^a \leq V_0$ and $V_0^b \leq V_0$. Moreover, for a newborn player to prefer to play a rather than b against a partner playing a , it must be that $V_0^b \leq V_0^a - (1 - \gamma)(u(b, a) - u(a, a))/\gamma$, which after substitutions implies $(1 - \mu^G)(u(b, a) - u(a, a)) \leq (1 - \mu^G)(\pi^G - \pi^B)$. When $\mu^G < 1$ (as is necessarily the case with noise), this inequality holds iff $u(b, a) - u(a, a) \leq \pi^G - \pi^B$. As $\mu^G \rightarrow 1$, this gives $u(b, a) - u(a, a) \leq u(a, a) - u(c, b)$, or Condition 4. Similarly, for a newborn player to prefer to play b rather than a against a partner playing c , it must be that $V_0^a \leq V_0^b + (1 - \gamma)(u(b, c) - u(a, c))/\gamma$, which implies $\mu^G(u(b, c) - u(a, c)) \geq (1 - \mu^G)(\pi^B - \pi^G)$.

As $\mu^G \rightarrow 1$, this gives $u(b, c) \geq u(a, c)$, or Condition 3. The actual proof relaxes the assumptions that newborn players are in good standing and have the highest continuation value, and instead shows that there exists some good-standing record at which an argument concerning incentive constraints similar to that given above implies Conditions 3 and 4.

6 Convergence

So far, we have focused on steady states without discussing stability or convergence properties. For general record systems and general strategies, we cannot say much about this. However, we now show that for the trigger strategies analyzed in Sections

4 and 5, there is always a unique steady state, and there is convergence to this steady state from an arbitrary initial distribution of records.¹⁷

Theorem 6. *Under the trigger strategies in Sections 4 and 5.1, there is a unique steady state, and the period- t state μ^t converges to the steady state from any initial state μ^0 .*

To see why this is the case, consider the trigger strategies in Section 5. For simplicity, consider just the shares corresponding to good-standing player scores. Since the probability that a given good-standing player increases their score in a given period is $(n - 2)\varepsilon$, these score shares in the population evolve according to

$$\begin{aligned}\mu_0^{t+1} &= 1 - \gamma + \gamma(1 - (n - 2)\varepsilon)\mu_0^t, \\ \mu_k^{t+1} &= \gamma(n - 2)\varepsilon\mu_{k-1}^t + \gamma(1 - (n - 2)\varepsilon)\mu_k^t \text{ for } 0 < k < K.\end{aligned}$$

Note that the evolution of μ_0^t depends only on its previous value, and since $\gamma(1 - (n - 2)\varepsilon) < 1$, μ_0^t must converge to a unique value, namely $(1 - \gamma)/(1 - \gamma + \gamma(n - 2)\varepsilon)$. Although the evolution of μ_1^t depends on the previous value of μ_0^t in addition to its own previous value, a similar argument shows that μ_1^t converge to a unique value because μ_0^t is approximately constant when t is high. This argument iterates to higher k so that each μ_k^t converges to a unique value. The proof is in A.6.

Likewise, there is also convergence to the unique steady state for the PPR trigger strategies in Section 5.2.

Theorem 7. *Under the PPR trigger strategies in Section 5.2, there is a unique steady state, and the period- t state μ^t converges to the steady state from any initial state μ^0 .*

Here the proof uses an equation governing the dynamics of the share of good-standing records to show that there must be a unique steady state, and shows that the

¹⁷This result depends on the uniformity of noise, which we assumed for the constructions in both sections. When noise is not uniform, there is not necessarily a unique steady state, but there are simple conditions on the initial population shares which guarantee convergence to the steady state with the highest share of good-standing players, and analogous conditions which guarantee convergence to the steady state with the lowest share of good-standing players.

value of $\mu^{t+1,G}$ is increasing in its previous value $\mu^{t,G}$, which results in convergence to the steady state from arbitrary initial population shares.

Finally, we note that for all of these strategy classes, the share of good-standing players remains arbitrarily close to its steady-state value over time if the initial population shares are sufficiently close to their steady-state values. Since the equilibria in this paper are all strict, it follows that they remain equilibria even when the population is not at steady-state, as long as the initial population shares are close enough to the steady state.

7 Discussion

This paper introduces a new environment for the study of repeated social interactions, where players interact with a sequence of anonymous and random opponents, and their limited information about opponents' past play (and the play of those their opponents have played, etc.) is summarized by "records." Unlike other papers with these features, our model does not have an initial time, supposes that players have geometrically distributed lifetimes, features noise, and emphasizes strict equilibria. We have focused here on the question of which outcomes can be fully supported by trigger strategies. Our headline results are that almost any outcome can be fully supported with second-order records, while with first-order records an outcome can be fully supported if it has a corresponding unprofitable punishment, but often cannot be supported otherwise.

We conclude by discussing some possible extensions and alternative models.

More complex strategies. This paper focuses on trigger strategies, and it would be interesting to consider more complex classes of strategies. In Clark, Fudenberg, and Wolitzky (2019), we do this for the prisoner's dilemma, where we give necessary and nearly-sufficient conditions for supporting cooperation in the class of all strategies that satisfy a "coordination-proofness" condition (a class that greatly generalizes trigger strategies). Performing a similar analysis for other games, or characterizing cooperation under non-coordination-proof strategies in the prisoner's dilemma, are left for future

work. We note however that there is no obvious single notion of “general strategies” in our model. For example, it can make a difference whether or not PPR are allowed, and also whether some type of symmetry-breaking device is allowed when two players with the same record meet. This ambiguity is another reason why we have focused on relatively simple strategies.

Sequential moves. The trigger strategies used to prove Theorem 2 specify that a good-standing player plays b against a bad-standing player (who is supposed to play c). This would not be a perfect equilibrium with sequential moves if the bad-standing player moves first: If a bad-standing player with a good-standing opponent could play a and the opponent observes this before taking their own action, the opponent would prefer to play a as well, and then the bad-standing player would receive $u(a, a)$, which exceeds the punishment payoff $u(c, b)$.

More generally, with any first-order or second-order record system, if any player can “jump the gun” when matched with an opponent with record r' , that player can guarantee a payoff of $\max_r u(\mathbf{s}(r, r'), \mathbf{s}(r', r))$ in this match. This implies that all players must receive the same payoff when matched with each possible opponent, which in turn implies that only stage game equilibrium behavior can be supported.

We have three responses to this observation. First, jumping the gun is often impossible: our model applies not only when actions are literally simultaneous, but also whenever both players must choose their actions before observing the opponent’s action, which seems like a natural reduced-form model for the often-realistic case where cooperative behavior unfolds gradually within each match. Second, all of our results go through with sequential moves when good-standing players move before bad-standing players. Thus, our results are unchanged when moves are sequential and the timing of moves can be specified as a part of equilibrium behavior. Finally, if one requires equilibria to be robust to all possible sequential move orders, then the trigger strategies used to prove Theorem 2 remain applicable if records are interdependent rather than second order. Thus, while second-order records are more informationally robust than interdependent records, interdependent records are more robust against manipulations

to the order of moves.

More general interaction structures. Our model assumes that each player is matched to play the game every period. The same model describes the steady states when some constant non-zero share of players is selected at random to play each period, with γ now interpreted as the probability of surviving for one more interaction. If different players are matched to play the game at different frequencies, the steady-state equations would be the same, but different players would face different incentive constraints. This extension seems interesting but we do not cover it here, except to note that, since our equilibria are strict, they are robust to small differences in interaction frequencies.

Multiple populations. It is easy to adapt our model to settings with multiple populations of players. One interesting observation here is that efficient outcomes can always be fully supported in situations with one-sided incentive problems.¹⁸ For example, suppose a population of player 1's and a population of player 2's repeatedly play the 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). Here 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 record is $k < K$, and both play D if player 1's record is $k \geq K$.

Endogenous records. This paper has considered how features of an exogenously given information structure (e.g. whether records are first-order, second-order, or interdependent) 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, as 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.¹⁹

¹⁸Proposition 4 of Kandori (1992) is a similar result in a fixed-population model without noise.

¹⁹Another possibility would be to consider "optimal" record systems subject to some constraints,

References

- [1] Dilip Abreu and Ariel Rubinstein. “The Structure of Nash Equilibrium in Repeated Games with Finite Automata”. *Econometrica* (1988).
- [2] Charalambos D. Aliprantis and Kim C. Border. *Infinite Dimensional Analysis: A Hitchhiker’s Guide 3rd Edition*. Third. Berlin: Springer, 2006.
- [3] Nicolas Baumard. “Has punishment played a role in the evolution of cooperation? A critical review”. *Mind & Society* (2010).
- [4] Lisa Bernstein. “Opting out of the legal system: Extralegal contractual relations in the diamond industry”. *The Journal of Legal Studies* (1992).
- [5] V Bhaskar, George J Mailath, and Stephen Morris. “Purification in the infinitely-repeated prisoners’ dilemma”. *Review of Economic Dynamics* (2008).
- [6] V Bhaskar and Caroline Thomas. “Community Enforcement of Trust with Bounded Memory”. *Review of Economic Studies* (2018).
- [7] Samuel Bowles and Herbert Gintis. *A cooperative species: Human reciprocity and its evolution*. Princeton, NJ, US. Princeton University Press, 2011.
- [8] Robert Boyd et al. “The evolution of altruistic punishment”. *Proceedings of the National Academy of Sciences* (2003).
- [9] Daniel Clark, Drew Fudenberg, and Alexander Wolitzky. *Steady-State Equilibria in Anonymous Repeated Games, II: Coordination-Proof Strategies in the Prisoner’s Dilemma*. Tech. rep. 2019.
- [10] Karen Clay. “Trade without law: private-order institutions in Mexican California”. *The Journal of Law, Economics, and Organization* (1997).
- [11] Olivier Compte and Andrew Postlewaite. “Plausible cooperation”. *Games and Economic Behavior* (2015).
- [12] Jacques Cremer. “Cooperation in Ongoing Organizations”. *The Quarterly Journal of Economics* (1986).
- [13] Pedro Dal Bó and Guillaume R Fréchette. “On the Determinants of Cooperation in Infinitely Repeated Games: A Survey”. *Journal of Economic Literature* (2018).
- [14] Joyee Deb, Takuo Sugaya, and Alexander Wolitzky. “The Folk Theorem in Repeated Games with Anonymous Random Matching”. 2018.
- [15] Chrysanthos Dellarocas. “Reputation mechanism design in online trading environments with pure moral hazard”. *Information systems research* (2005).
- [16] Robert C Ellickson. *Order without law*. Harvard University Press, 1991.

in the spirit of information design.

- [17] Glenn Ellison. “Cooperation in the prisoner’s dilemma with anonymous random matching”. *Review of Economic Studies* (1994).
- [18] James D Fearon and David D Laitin. “Explaining interethnic cooperation”. *American political science review* (1996).
- [19] Eric J Friedman and Paul Resnick. “The social cost of cheap pseudonyms”. *Journal of Economics & Management Strategy* (2001).
- [20] James W. Friedman. “A Non-cooperative Equilibrium for Supergames”. *The Review of Economic Studies* (1971).
- [21] Drew Fudenberg and Kevin He. “Learning and type compatibility in signaling games”. *Econometrica* (2018).
- [22] Drew Fudenberg and Kevin He. “Player-Compatible Equilibrium”. 2019.
- [23] Drew Fudenberg, David Levine, and Eric Maskin. “The Folk Theorem with Imperfect Public Information”. *Econometrica* (1994).
- [24] Drew Fudenberg, David G. Rand, and Anna Dreber. “Slow to Anger and Fast to Forgive: Cooperation in an Uncertain World”. *American Economic Review* (2012).
- [25] Herbert Gintis et al. “Explaining altruistic behavior in humans”. *Evolution and human Behavior* (2003).
- [26] Avner Greif. “Contract enforceability and economic institutions in early trade: The Maghribi traders’ coalition”. *The American economic review* (1993).
- [27] Avner Greif, Paul Milgrom, and Barry R Weingast. “Coordination, commitment, and enforcement: The case of the merchant guild”. *Journal of political economy* (1994).
- [28] Francesco Guala. “Reciprocity: Weak or strong? What punishment experiments do (and do not) demonstrate”. *Behavioral and Brain Sciences* (1 2012).
- [29] Yuval Heller. “Instability of belief-free equilibria”. *Journal of Economic Theory* (2017).
- [30] Yuval Heller and Erik Mohlin. “Observations on Cooperation”. *Review of Economic Studies* (2018).
- [31] Yongjoon Joe et al. “Automated Equilibrium Analysis of Repeated Games with Private Monitoring: A POMDP Approach”. *Conference on Autonomous Agents and Multiagent Systems*. 2012.
- [32] Michihiro Kandori. “Social Norms and Community Enforcement”. *The Review of Economic Studies* (1992).
- [33] Benjamin Klein and Keith B Leffler. “The role of market forces in assuring contractual performance”. *Journal of political Economy* (1981).

- [34] Daniel B Klein. “Promise keeping in the great society: A model of credit information sharing”. *Economics & Politics* (1992).
- [35] 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* (1990).
- [36] Markus M Möbius. “Trading Favors”. 2001.
- [37] Martin A. Nowak and Karl Sigmund. “Evolution of indirect reciprocity by image scoring”. *Nature* (1998).
- [38] Masahiro Okuno-Fujiwara and Andrew Postlewaite. “Social Norms and Random Matching Games”. *Games and Economic Behavior* (1995).
- [39] Elinor Ostrom. *Governing the commons: The evolution of institutions for collective action*. Cambridge university press, 1990.
- [40] A Jorge Padilla and Marco Pagano. “Sharing default information as a borrower discipline device”. *European Economic Review* (2000).
- [41] Karthik Panchanathan and Robert Boyd. “A tale of two defectors: the importance of standing for evolution of indirect reciprocity”. *Journal of Theoretical Biology* (2003).
- [42] 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.
- [43] Robert W. Rosenthal. “Sequences of Games with Varying Opponents”. *Econometrica* (1979).
- [44] Robert W. Rosenthal and Henry J. Landau. “A game-theoretic analysis of bargaining with reputations”. *Journal of Mathematical Psychology* (1979).
- [45] Ariel Rubinstein. “Finite automata play the repeated prisoner’s dilemma”. *Journal of Economic Theory* (1986).
- [46] David J. Salant. “A repeated game with finitely lived overlapping generations of players”. *Games and Economic Behavior* (1991).
- [47] Lones Smith. “Folk Theorems in Overlapping Generations Games”. *Games and Economic Behavior* (1992).
- [48] Satoru Takahashi. “Community enforcement when players observe partners’ past play”. *Journal of Economic Theory* (2010).

Appendix

A.1 Proof of Theorem 1

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 5. *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 $\lambda > (1 - \eta)/(1 - \eta/2)$. 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 larger than λ , i.e. $\sum_{r'' \in S(r, m, a, a')} \rho(r, r', a, a')[r''] > \lambda$ 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. Since records and exits are independent, the probability that a 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 $\sum_{r \in L(t, \eta)} f_{\sigma}(\mu)[r] > (1 - \eta)(1 - \gamma)\gamma^t$. ■

The next corollary is an immediate consequence of Properties 2 and 3 of Lemma 5.

Corollary 4. *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 of finite subsets of R , $\{L(t, \eta)\}_{t \in \mathbb{N}, \eta > 0}$, satisfying the three properties in Lemma 5 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 4, must contain every steady-state distribution μ . The next lemma uses Corollary 4 to show that \bar{M} is non-empty.

Lemma 6. *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 $t \in \mathbb{N}$, we have by Corollary 4 that

$$\sum_{r \in L(t, \eta)} \mu_r^{i_j} \geq (1 - \eta)(1 - \gamma)\gamma^t$$

for all $\eta > 0$ for all sufficiently high $j \in \mathbb{N}$, so

$$\sum_{r \in L(t, \eta)} \tilde{\mu}_r \geq (1 - \eta)(1 - \gamma)\gamma^t. \tag{3}$$

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 (3) 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 (3) implies that $\tilde{\mu} \in \bar{M}$. ■

The following three claims imply that f_σ has a fixed point in \bar{M} .²⁰

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 6 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, $|\mu_r^{i_j} - \tilde{\mu}_r| < \eta/2$ for all $r \notin L(\eta/2)$ for all $j \in \mathbb{N}$. Now, let $J \in \mathbb{N}$ be such that $|\mu_r^{i_j} - \tilde{\mu}_r| < \eta/2$ 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 $j > J$. ■

Proof of Claim 2. For any $\mu \in \bar{M}$, Properties 2 and 3 of Lemma 5 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 $\sum_{r \in R(t)} f_\sigma(\mu)[r] = (1 - \gamma)\gamma^t$. ■

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 a given $\eta > 0$, 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 focus on the case where $r \neq 0$.

²⁰This follows from Corollary 17.56 (page 583) of Aliprantis and Border (2006), 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 show that no steady state exists for the interdependent record system with $R = \mathbb{N}$ and $\rho(r, r') = \max\{r, r'\} + 1$, whenever $\gamma > 1/2$. 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 we must have $\sum_{r=r^*}^{\infty} f(\mu)[r] = \mu_*$, a contradiction.

A.2 Proof of Lemma 1 and Theorem 2

A.2.1 Proof of Lemma 1

The feasibility constraint of Lemma 1 comes from the following lemma, the proof of which is in OA.1.1.

Lemma 7. *In a trigger strategy equilibrium with tolerance K and share of good-standing players μ^G , $\mu_k = \alpha(\gamma, \varepsilon)^k(1 - \alpha(\gamma, \varepsilon))$ for $k < K$.*

To see why the feasibility constraint of Lemma 1 comes from Lemma 7, note that

$$\mu^G = \sum_{k=0}^{K-1} \alpha(\gamma, \varepsilon)^k(1 - \alpha(\gamma, \varepsilon)) = 1 - \alpha(\gamma, \varepsilon)^K.$$

We turn to the incentives constraint in Lemma 1. The next lemma gives closed forms for the equilibrium values as a function of the player's score.

Lemma 8. *In a trigger strategy equilibrium with tolerance K and share of good-standing players μ^G ,*

$$V_k = (1 - \alpha(\gamma, \varepsilon)^{K-k})(\mu^G u(a, a) + (1 - \mu^G)u(b, c)) + \alpha(\gamma, \varepsilon)^{K-k}(\mu^G u(c, b) + (1 - \mu^G)u(d, d))$$

for $k < K$, and

$$V_k = \mu^G u(c, b) + (1 - \mu^G)u(d, d)$$

for $k \geq K$.

The proof of this lemma, which is simple algebra, is in OA.1.2.

To make use of this result, note that the payoff to a good-standing player with score k from playing a against an opponent playing a is $(1 - \gamma)u(a, a) + \gamma(1 - \nu(\varepsilon))V_k + \gamma\nu(\varepsilon)V_{k+1}$, while their payoff from playing some action x is $(1 - \gamma)u(x, a) + \gamma(1 - p_x(\varepsilon))V_k + \gamma p_x(\varepsilon)V_{k+1}$, where $p_x(\varepsilon)$ gives the probability that the player's score increases when they play x . Thus, a good-standing player with score k prefers to play a against a rather than x against a iff

$$(1 - \nu(\varepsilon) - p_x(\varepsilon))\frac{\gamma}{1 - \gamma}(V_k - V_{k+1}) > u(x, a) - u(a, a).$$

Using the expressions for the V_k in Lemma 8 gives

$$\frac{1 - \nu(\varepsilon) - p_x(\varepsilon)}{\nu(\varepsilon)}\alpha(\gamma, \varepsilon)^{K-k}(\mu^G(u(a, a) - u(c, b)) + (1 - \mu^G)(u(b, c) - u(d, d))) > u(x, a) - u(a, a).$$

Note that the left hand side of this inequality is increasing in k , so a necessary and sufficient condition for all good-standing players to play a rather than any x against a is the version of this inequality for $k = 0$, which using the fact that $\mu^G = 1 - \alpha(\gamma, \varepsilon)^K$, is equivalent to

$$\frac{1 - \nu(\varepsilon) - p_x(\varepsilon)}{\nu(\varepsilon)}(1 - \mu^G)(\mu^G(u(a, a) - u(c, b)) + (1 - \mu^G)(u(b, c) - u(d, d))) > \max_x u(x, a) - u(a, c).$$

Since $1 - \nu(\varepsilon) - p_x(\varepsilon)$ is bounded above 0, it follows that there is some z such that whenever

$$\frac{1}{\nu(\varepsilon)}(1 - \mu^G)(\mu^G(u(a, a) - u(c, b)) + (1 - \mu^G)(u(b, c) - u(d, d))) > z,$$

it is optimal for good-standing players to prefer to play a against a . A similar argument shows that for sufficiently high z , this inequality implies that it is optimal for good-standing players to prefer to play b against c .

A.2.2 Proof of Theorem 2

Proof of Theorem 2. Fix $\eta \in (0, 1)$. For arbitrary (γ, ε) , consider the trigger strategy with tolerance $K(\gamma, \varepsilon)$ set by

$$K(\gamma, \varepsilon) = \left\lceil \frac{\ln(\eta)}{\ln(\alpha(\gamma, \varepsilon))} \right\rceil,$$

so that $K(\gamma, \varepsilon)$ is the smallest integer greater than or equal to $\ln(\eta)/\ln(\alpha(\gamma, \varepsilon))$. Note that with this tolerance the corresponding share of good-standing players is $\mu^G(\gamma, \varepsilon) = 1 - \alpha(\gamma, \varepsilon)^{K(\gamma, \varepsilon)}$. By construction, $\mu^G(\gamma, \varepsilon) \geq 1 - \eta$ for all (γ, ε) , and, whenever $K(\gamma, \varepsilon) > 1$, $\mu^G(\gamma, \varepsilon) \leq 1 - \eta^2$.

We will now argue that for small η , whenever γ is sufficiently large and ε is sufficiently small, such a strategy (along with the corresponding steady state) constitutes an equilibrium. In particular, suppose that η is such that the difference in flow payoff between a good-standing player and a bad-standing player is strictly positive whenever $\mu^G \geq 1 - \eta$. That is, $(1 - \eta)(u(a, a) - u(c, b)) + \eta(u(b, c) - u(d, d)) > 0$.

Consider first the case where $K(\gamma, \varepsilon) > 1$. Since $\lim_{\varepsilon \rightarrow 0} 1/\nu(\varepsilon) = \infty$ and $1 - \eta \leq \mu^G(\gamma, \varepsilon) \leq 1 - \eta^2$ whenever $K(\gamma, \varepsilon) > 1$, it follows that for ε sufficiently small, this trigger strategy is an equilibrium whenever $K(\gamma, \varepsilon) > 1$.

Now consider the case where $K(\gamma, \varepsilon) = 1$. In this case, $\mu^G = 1 - \alpha(\gamma, \varepsilon)$, so

$$\frac{1}{\nu(\varepsilon)}(1 - \mu^G(\gamma, \varepsilon)) = \frac{\gamma}{1 - \gamma + \gamma\nu(\varepsilon)}.$$

Since $\lim_{(\gamma, \varepsilon) \rightarrow (1, 0)} \gamma/(1 - \gamma + \gamma\nu(\varepsilon)) = \infty$, it follows that for γ sufficiently large and ε sufficiently small, this trigger strategy is an equilibrium whenever $K(\gamma, \varepsilon) = 1$. ■

A.3 Proof of Lemma 2 and Theorem 3

A.3.1 Proof of Lemma 2

The feasibility constraint of Lemma 2 comes from the following lemma, which is proved in OA.2.1.

Lemma 9. *In a trigger strategy equilibrium with tolerance K and share of good-standing players μ^G , for all $k < K$, $\mu_k = \beta(\gamma, \varepsilon)^k(1 - \beta(\gamma, \varepsilon))$.*

To see why the feasibility constraint of Lemma 2 comes from Lemma 9, note that

$$\mu^G = \sum_{k=0}^{K-1} \beta(\gamma, \varepsilon)^k(1 - \beta(\gamma, \varepsilon)) = 1 - \beta(\gamma, \varepsilon)^K.$$

We turn to the incentives conditions in Lemma 1. The next lemma gives closed forms for the equilibrium values as a function of the player's score.

Lemma 10. *In a trigger strategy equilibrium with tolerance K and share of good-standing players μ^G ,*

$$V_k = (1 - \beta(\gamma, \varepsilon)^{K-k})(\mu^G u(a, a) + (1 - \mu^G)u(b, c)) + \beta(\gamma, \varepsilon)^{K-k}(\mu^G u(c, b) + (1 - \mu^G)u(d, d))$$

for $k < K$, and

$$V_k = \mu^G u(c, b) + (1 - \mu^G)u(d, d)$$

for $k \geq K$.

The proof of this lemma, which is simple algebra, is in OA.2.2.

To make use of this result, note that the payoff to a good-standing player with score k from playing a against an opponent playing a is $(1 - \gamma)u(a, a) + \gamma(1 - (n - 2)\varepsilon)V_k + \gamma(n - 2)\varepsilon V_{k+1}$, while their payoff from playing some action $x \neq a, b$ is $(1 - \gamma)u(x, a) + \gamma 2\varepsilon V_k + \gamma(1 - 2\varepsilon)V_{k+1}$.²¹ Thus, a good-standing player with score k prefers to play a

²¹Because $u(a, a) > u(b, a)$, we do not have to worry about a good-standing player playing b rather than a against an opponent playing a .

against a rather than $x \neq a, b$ against a iff

$$(1 - n\varepsilon) \frac{\gamma}{1 - \gamma} (V_k - V_{k+1}) > u(x, a) - u(a, a).$$

Using the expressions for the V_k in Lemma 10 gives

$$\frac{1 - n\varepsilon}{(n - 2)\varepsilon} \beta(\gamma, \varepsilon)^{K-k} (\mu^G(u(a, a) - u(c, b)) + (1 - \mu^G)(u(b, c) - u(d, d))) > u(x, a) - u(a, a)$$

Note that the left hand side of this inequality is increasing in k , so a necessary and sufficient condition for all good-standing players to play a rather than any $x \neq a, b$ against a is the version of this inequality for $k = 0$, which using the fact that $\mu^G = 1 - \beta(\gamma, \varepsilon)^K$, is equivalent to

$$\frac{1 - n\varepsilon}{(n - 2)\varepsilon} (1 - \mu^G)(\mu^G(u(a, a) - u(c, b)) + (1 - \mu^G)(u(b, c) - u(d, d))) > \max_x u(x, a) - u(a, c).$$

It can be similarly shown that a necessary and sufficient condition for all good-standing players to prefer to play b rather than any $x \neq a, b$ against c is

$$\frac{1 - n\varepsilon}{(n - 2)\varepsilon} (1 - \mu^G)(\mu^G(u(a, a) - u(c, b)) + (1 - \mu^G)(u(b, c) - u(d, d))) > \max_x u(x, c) - u(b, c).$$

These two conditions together give the incentive constraint in Lemma 2.

A.3.2 Proof of Theorem 3

Proof of Theorem 3. Fix $\eta \in (0, 1)$. For arbitrary (γ, ε) , consider the trigger strategy with tolerance $K(\gamma, \varepsilon)$ set by

$$K(\gamma, \varepsilon) = \left\lceil \frac{\ln(\eta)}{\ln(\beta(\gamma, \varepsilon))} \right\rceil,$$

so that $K(\gamma, \varepsilon)$ is the smallest integer greater than or equal to $\ln(\eta)/\ln(\beta(\gamma, \varepsilon))$. Note that with this tolerance the corresponding share of good-standing players is $\mu^G(\gamma, \varepsilon) =$

$1 - \beta(\gamma, \varepsilon)^{K(\gamma, \varepsilon)}$. By construction, $\mu^G(\gamma, \varepsilon) \geq 1 - \eta$ for all (γ, ε) , and, whenever $K(\gamma, \varepsilon) > 1$, $\mu^G(\gamma, \varepsilon) \leq 1 - \eta^2$.

We will now argue that for small η , whenever γ is sufficiently large and ε is sufficiently small, such a strategy constitutes an equilibrium. In particular, suppose that η is such that the difference in flow payoff between a good-standing player and a bad-standing player is strictly positive whenever $\mu^G \geq 1 - \eta$. That is, $(1 - \eta)(u(a, a) - u(c, b)) + \eta(u(b, c) - u(d, d)) > 0$.

Consider first the case where $K(\gamma, \varepsilon) > 1$. Since $\lim_{\varepsilon \rightarrow 0} (1 - n\varepsilon)/((n - 2)\varepsilon) = \infty$ and $1 - \eta \leq \mu^G(\gamma, \varepsilon) \leq 1 - \eta^2$ whenever $K(\gamma, \varepsilon) > 1$, it follows that for ε sufficiently small, this trigger strategy is an equilibrium whenever $K(\gamma, \varepsilon) > 1$.

Now consider the case where $K(\gamma, \varepsilon) = 1$. In this case, $\mu^G = 1 - \beta(\gamma, \varepsilon)$, so

$$\frac{1 - n\varepsilon}{(n - 2)\varepsilon}(1 - \mu^G(\gamma, \varepsilon)) = \frac{\gamma(1 - n\varepsilon)}{1 - \gamma + \gamma(n - 2)\varepsilon}.$$

Since $\lim_{(\gamma, \varepsilon) \rightarrow (1, 0)} \gamma(1 - n\varepsilon)/(1 - \gamma + \gamma(n - 2)\varepsilon) = \infty$, it follows that for γ sufficiently large and ε sufficiently small, this trigger strategy is an equilibrium whenever $K(\gamma, \varepsilon) = 1$. ■

A.4 Proof of Theorem 4

Let $h_n : (0, 1)^5 \rightarrow \mathbb{R}$ be the function given by

$$h_n(\gamma, \varepsilon, \mu^G, \tau_a, \chi) = \frac{1 - \mu^G - (n - 2)\mu^G \varepsilon^{\frac{\gamma}{1-\gamma}} \chi - (1 - (n - 2)\varepsilon)\mu^G \frac{\gamma}{1-\gamma} \tau_a}{(\varepsilon \mu^G + (1 - \mu^G)(1 - (n - 1)\varepsilon))\mu^G}. \quad (4)$$

Lemma 11. *There is an equilibrium with share μ^G of players in G iff the feasibility constraint*

$$\tau_b - \tau_a = \frac{1 - \gamma}{\gamma} h_n(\gamma, \varepsilon, \mu^G, \tau_a, \chi),$$

and incentive constraints

$$(a \text{ vs } b|a) : (1 - n\varepsilon)h_n(\gamma, \varepsilon, \mu^G, \tau_a, \chi)(\mu^G(u(a, a) - u(c, b)) + (1 - \mu^G)(u(b, c) - u(d, d))) \\ > u(b, a) - u(a, a),$$

$$(b \text{ vs } a|c) : (1 - n\varepsilon)h_n(\gamma, \varepsilon, \mu^G, \tau_a, \chi)(\mu^G(u(a, a) - u(c, b)) + (1 - \mu^G)(u(b, c) - u(d, d))) \\ < u(b, c) - u(a, c),$$

$$(a \text{ vs } x|a) : (1 - n\varepsilon)\frac{\gamma}{1 - \gamma}(\chi - \tau_a)\mu^G(\mu^G(u(a, a) - u(c, b)) + (1 - \mu^G)(u(b, c) - u(d, d))) \\ > \max_{x \neq a, b} u(x, a) - u(a, a),$$

$$(b \text{ vs } x|c) : (1 - n\varepsilon)\frac{\gamma}{1 - \gamma}(\chi - \tau_b)\mu^G(\mu^G(u(a, a) - u(c, b)) + (1 - \mu^G)(u(b, c) - u(d, d))) \\ > \max_{x \neq a, b} u(x, c) - u(b, c),$$

are satisfied.

The proof of Lemma 11 is in OA.3. The feasibility constraints come from calculating the relationship between the transition probabilities τ_a , τ_b , and χ necessary to support the steady-state share of μ^G in G . The incentive constraints come from solving V^G and V^B (the value function of players in G and B), and using the transition probabilities to evaluate the good-standing players' incentives to follow the prescribed play.

Let $\kappa \in (0, 1)$ be such that

$$\frac{u(b, a) - u(a, a)}{u(a, a) - u(c, b)} < \kappa < \min \left\{ \frac{u(b, c) - u(a, c)}{u(a, a) - u(c, b)}, 1 \right\}$$

Fix $\mu^G \in (0, 1)$, and let $\chi(\gamma, \varepsilon)$ be given by

$$\chi(\gamma, \varepsilon) = \min \left\{ \frac{1 - \gamma(1 - \kappa)(1 - \mu^G)}{\gamma\varepsilon(n - 2)\mu^G}, 1 \right\},$$

$\tau_a(\gamma, \varepsilon)$ be given by

$$\tau_a(\gamma, \varepsilon) = \frac{1 - \gamma}{\gamma(1 - (n - 2)\varepsilon)\mu^G} \left((1 - \kappa)(1 - \mu^G) - (n - 2)\mu^G\varepsilon\frac{\gamma}{1 - \gamma}\chi(\gamma, \varepsilon) \right),$$

and $\tau_b(\gamma, \varepsilon)$ be given by

$$\tau_b(\gamma, \varepsilon) = \tau_a(\gamma, \varepsilon) + \frac{1-\gamma}{\gamma} h_n(\gamma, \varepsilon, \mu^G, \tau_a(\gamma, \varepsilon), \chi(\gamma, \varepsilon)).$$

Note that $\chi(\gamma, \varepsilon), \tau_a(\gamma, \varepsilon), \tau_b(\gamma, \varepsilon) \geq 0$ for all (γ, ε) , $\chi(\gamma, \varepsilon) \leq 1$ for all (γ, ε) , $\lim_{(\gamma, \varepsilon) \rightarrow (1, 0)} \tau_a(\gamma, \varepsilon) = 0$, and $\lim_{(\gamma, \varepsilon) \rightarrow (1, 0)} \tau_b(\gamma, \varepsilon) = 0$.

Lemma 12. *For μ^G close enough to 1, this constitutes an equilibrium for (γ, ε) sufficiently close to $(1, 0)$.*

Proof. We have already established that the feasibility constraints are satisfied for (γ, ε) sufficiently close to 1. Thus, we need only show that the incentive constraints are satisfied.

Note that

$$\lim_{(\gamma, \varepsilon) \rightarrow (1, 0)} (1 - n\varepsilon) h_n(\gamma, \varepsilon, \mu^G, \tau_a(\gamma, \varepsilon), \chi(\gamma, \varepsilon)) = \kappa.$$

Since

$$u(b, a) - u(a, a) < \kappa(u(a, a) - u(c, b)) < u(b, c) - u(a, c),$$

it thus follows that the $(a \text{ vs } b|a)$ constraint and the $(b \text{ vs } a|c)$ constraint are both satisfied in the $(\gamma, \varepsilon) \rightarrow (1, 0)$ limit when μ^G is sufficiently close to 1.

Additionally, note that

$$\lim_{(\gamma, \varepsilon) \rightarrow (1, 0)} (1 - n\varepsilon) \frac{\gamma}{1 - \gamma} (\chi(\gamma, \varepsilon) - \tau_a(\gamma, \varepsilon)) = \lim_{(\gamma, \varepsilon) \rightarrow (1, 0)} (1 - n\varepsilon) \frac{\gamma}{1 - \gamma} (\chi(\gamma, \varepsilon) - \tau_b(\gamma, \varepsilon)) = \infty.$$

Since $u(a, a) - u(c, b) > 0$, it follows that both the $(a \text{ vs } b|a)$ and $(b \text{ vs } a|c)$ constraints are satisfied in the $(\gamma, \varepsilon) \rightarrow (1, 0)$ limit when μ^G is sufficiently close to 1. \blacksquare

A.5 Proof of Theorem 5

Throughout this section, we assume that (a, a) is not a stage-game Nash equilibrium, and we restrict attention to the class of trigger strategies in which good-standing players

play a against fellow good-standing players and action b against bad-standing players, and bad-standing players play c against good-standing players and d against bad-standing players. We denote the expected continuation payoff of a record r player who plays action x by V_r^x and denote the probability that a record r player plays x by p_r^x . All omitted proofs are in OA.4.

Fix an equilibrium and 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 $\mu^G \pi^G + (1 - \mu^G) \pi^B$ (the average flow payoff in the population), we have $\bar{V} \geq V_0 = \mu^G \pi^G + (1 - \mu^G) \pi^B$.

Lemma 13. *For μ^G sufficiently close to 1, there is no sequence of bad-standing records $\{r_n\}_{n \in \mathbb{N}}$ such that $\lim_{n \rightarrow \infty} V_{r_n} = \bar{V}$.*

Lemma 14. *If for every $\eta > 0$ there exist some γ and ε such that trigger strategies support share at least $1 - \eta$ of the population playing a , then $u(a, a) \geq u(c, b)$.*

Lemma 15. *If $\mu^G > 0$, then $u(a, a) - u(b, a) > u(a, c) - u(b, c)$.*

Proof. For a good-standing player with record r to prefer to play a against a than b against a , it must be that $(1 - \gamma)u(a, a) + \gamma V_r^a > (1 - \gamma)u(b, a) + \gamma V_r^b$, which implies $u(a, a) - u(b, a) > \frac{\gamma}{1 - \gamma}(V_r^b - V_r^a)$. Likewise, we can show that for a good-standing player with record r to prefer to play b against c rather than a against c , it must be that $u(a, c) - u(b, c) < \frac{\gamma}{1 - \gamma}(V_r^b - V_r^a)$. Combining these two inequalities gives $u(a, a) - u(b, a) > u(a, c) - u(b, c)$. \blacksquare

Let $D_r^{a,b} := \gamma(V_r^a - V_r^b)/(1 - \gamma)$.

Lemma 16. $D_r^{a,b} = \frac{1}{1 - \mu^G} \left(\pi^G - V_r - \frac{\gamma}{1 - \gamma}(V_r - V_r^a) \right) = -\frac{1}{\mu^G} \left(\pi^G - V_r - \frac{\gamma}{1 - \gamma}(V_r - V_r^b) \right)$.

Proof. This follows immediately from $V_r = (1 - \gamma)\pi^G + \gamma\mu^G V_r^a + \gamma(1 - \mu^G)V_r^b$. \blacksquare

Lemma 17. *For μ^G sufficiently close to 1, there is some good-standing record r' such that*

$$\min \left\{ V_{r'} - \frac{\gamma}{1 - \gamma}(V_{r'}^a - V_{r'}), V_{r'} - \frac{\gamma}{1 - \gamma}(V_{r'}^b - V_{r'}) \right\} \geq \mu^G \pi^G + (1 - \mu^G) \pi^B. \quad (5)$$

Proof of Lemma 17. First consider the case where $\bar{V} = \mu^G \pi^G + (1 - \mu^G) \pi^B$. Then there must be some record k' such that $V_{k'} = \mu^G \pi^G + (1 - \mu^G) \pi^B$. By Lemma 13, such a k' cannot be a bad-standing record. Moreover $V_{k'}^a \leq \bar{V}$, so $V_{k'} - (\gamma/(1 - \gamma))(V_{k'}^a - V_{k'}) \geq \mu^G \pi^G + (1 - \mu^G) \pi^B$. Likewise, $V_{k'}^b \leq \bar{V}$, so $V_{k'} - (\gamma/(1 - \gamma))(V_{k'}^b - V_{k'}) \geq \mu^G \pi^G + (1 - \mu^G) \pi^B$.

Now, consider the case where $\bar{V} > \mu^G \pi^G + (1 - \mu^G) \pi^B$. For any sequence of records $\{r_n\}_{n \in \mathbb{N}}$ such that $\lim_{n \rightarrow \infty} V_{r_n} = \bar{V}$, $\lim_{n \rightarrow \infty} \max\{V_r^a - V_r, V_r^b - V_r, 0\} = 0$, so for all sufficiently high n we have $\max\{V_r^a - V_r, V_r^b - V_r\} \geq \mu^G \pi^G + (1 - \mu^G) \pi^B$. Additionally, by Lemma 13, for sufficiently high n , the record r_n must be a good-standing record. ■

Lemma 18. *Consider a good-standing record r' whose value and continuation values satisfy the inequalities in (5). As $\mu^G \rightarrow 1$, $D_{r'}^{a,b} < u(b, c) - u(a, c)$ implies $u(b, c) \geq u(a, c)$ and $D_{r'}^{a,b} > u(b, a) - u(a, a)$ implies $u(b, a) - u(a, a) \leq u(a, a) - u(c, b)$.*

Proof. Combining (5) with Lemma 16 and $D_{r'}^{a,b} < u(b, c) - u(a, c)$ gives

$$-\frac{1 - \mu^G}{\mu^G} (\pi^G - \pi^B) < u(b, c) - u(a, c).$$

As $\mu^G \rightarrow 1$, the left-hand side of this inequality converges to 0, which gives $u(b, c) \geq u(a, c)$, or Condition 3. Likewise, Lemma 16 and $D_{r'}^{a,b} > u(b, a) - u(a, a)$ requires that $\pi^G - \pi^B > u(b, a) - u(a, a)$. As $\mu^G \rightarrow 1$, the left-hand side of this inequality converges to $u(a, a) - u(b, a)$, which gives $u(b, a) - u(a, a) \leq u(a, a) - u(c, b)$, or Condition 4. ■

A.6 Proof of Theorem 6

Fix parameters $\gamma, \varepsilon \in (0, 1)$ and suppose that all players play according to the trigger strategy with tolerance level K . Given the value of μ_k^t , the population shares at the various scores, at time t , the corresponding value of the shares at time $t + 1$ are given

by

$$\begin{aligned}
\mu_0^{t+1} &= 1 - \gamma + \gamma(1 - (n-2)\varepsilon)\mu_0^t, \\
\mu_k^{t+1} &= \gamma(n-2)\varepsilon\mu_{k-1}^t + \gamma(1 - (n-2)\varepsilon)\mu_k^t \text{ for } 0 < k < K, \\
\mu_K^{t+1} &= \gamma(n-2)\varepsilon\mu_{K-1}^t + \gamma 2\varepsilon\mu_K^t, \\
\mu_k^{t+1} &= \gamma(1 - 2\varepsilon)\mu_{k-1}^t + \gamma 2\varepsilon\mu_k^t.
\end{aligned}$$

Next we show that $\lim_{t \rightarrow \infty} \mu_k^t = \beta(\gamma, \varepsilon)^k(1 - \beta(\gamma, \varepsilon))$ for all $k < K$ by induction over k . A similar argument can then be used to show that

$$\lim_{t \rightarrow \infty} \mu_k^t = \beta(\gamma, \varepsilon)^K \left(\frac{\gamma(1 - 2\varepsilon)}{1 - \gamma 2\varepsilon} \right)^{k-K} \frac{1 - \gamma}{1 - \gamma 2\varepsilon}$$

for all $k \geq K$.

Note that $\mu_0^{t+1} = 1 - \gamma + \gamma(1 - (n-2)\varepsilon)\mu_0^t$ is equivalent to $\mu_0^{t+1} - (1 - \beta(\gamma, \varepsilon)) = \gamma(1 - (n-2)\varepsilon)(\mu_0^t - (1 - \beta(\gamma, \varepsilon)))$. Thus $\mu_0^t = 1 - \beta(\gamma, \varepsilon) + (\gamma(1 - (n-2)\varepsilon))^t(\mu_0^0 - (1 - \beta(\gamma, \varepsilon)))$, so $\lim_{t \rightarrow \infty} \mu_0^t = 1 - \beta(\gamma, \varepsilon)$ since $\lim_{t \rightarrow \infty} (\gamma(1 - (n-2)\varepsilon))^t = 0$.

Suppose that $\lim_{t \rightarrow \infty} \mu_{k-1}^t = \beta(\gamma, \varepsilon)^{k-1}(1 - \beta(\gamma, \varepsilon))$. Note that $\mu_k^{t+1} = \gamma(n-2)\varepsilon\mu_{k-1}^t + \gamma(1 - (n-2)\varepsilon)\mu_k^t$ is equivalent to

$$\mu_k^{t+1} - \beta(\gamma, \varepsilon)\mu_{k-1}^t = \gamma(1 - (n-2)\varepsilon)(\mu_k^t - \beta(\gamma, \varepsilon)\mu_{k-1}^{t-1}) + \gamma(1 - (n-2)\varepsilon)\beta(\gamma, \varepsilon)(\mu_{k-1}^{t-1} - \mu_{k-1}^t),$$

which implies

$$\begin{aligned}
\mu_k^{t+1} - \beta(\gamma, \varepsilon)\mu_{k-1}^t &= \gamma(1 - (n-2)\varepsilon)^t(\mu_k^t - \beta(\gamma, \varepsilon)\mu_{k-1}^0) \\
&\quad + \sum_{i=1}^t \gamma(1 - (n-2)\varepsilon)^i \beta(\gamma, \varepsilon)(\mu_{k-1}^{t-i} - \mu_{k-1}^{t-i+1}).
\end{aligned}$$

This gives $\lim_{t \rightarrow \infty} \mu_k^t = \beta(\gamma, \varepsilon)^k(1 - \beta(\gamma, \varepsilon))$ when $\lim_{t \rightarrow \infty} \mu_{k-1}^t = \beta(\gamma, \varepsilon)^{k-1}(1 - \beta(\gamma, \varepsilon))$.

Online Appendix for “Steady-State Equilibria in Anonymous Repeated Games, I: Trigger-Strategy Equilibria in General Stage Games”

Daniel Clark, Drew Fudenberg, and Alexander Wolitzky

August 30, 2019

OA.1 Proof of Supporting Results for Lemma 1

OA.1.1 Proof of Lemma 7

Lemma 7. *In a trigger strategy equilibrium with tolerance K and share of good-standing players μ^G , $\mu_k = \alpha(\gamma, \varepsilon)^k(1 - \alpha(\gamma, \varepsilon))$ for $k < K$.*

Proof. The inflow into score 0 is $1 - \gamma$, while the outflow from score 0 is $(1 - \gamma)\mu_0 + \gamma\nu(\varepsilon)\mu^0$. Setting these equal gives

$$\mu_0 = \frac{1 - \gamma}{1 - \gamma + \gamma\nu(\varepsilon)} = 1 - \alpha(\gamma, \varepsilon).$$

Additionally, for every $0 < k < K$, both score k and score $k - 1$ are good standing scores. Thus, the inflow into score k is $\gamma\nu(\varepsilon)\mu_{k-1}$, while the outflow from score k is

*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.

$(1 - \gamma)\mu_k + \gamma\nu(\varepsilon)\mu_k$. Setting these equal gives

$$\mu_k = \frac{\gamma\nu(\varepsilon)}{1 - \gamma + \gamma\nu(\varepsilon)}\mu_{k-1} = \alpha(\gamma, \varepsilon)\mu_{k-1}.$$

Combining this with $\mu_0 = 1 - \alpha(\gamma, \varepsilon)$ gives $\mu_k = \alpha(\gamma, \varepsilon)^k(1 - \alpha(\gamma, \varepsilon))$ for $0 \leq k \leq K - 1$. ■

OA.1.2 Proof of Lemma 8

Lemma 8. *In a trigger strategy equilibrium with tolerance K and share of good-standing players μ^G ,*

$$V_k = (1 - \alpha(\gamma, \varepsilon)^{K-k})(\mu^G u(a, a) + (1 - \mu^G)u(b, c)) + \alpha(\gamma, \varepsilon)^{K-k}(\mu^G u(c, b) + (1 - \mu^G)u(d, d))$$

for $k < K$, and

$$V_k = \mu^G u(c, b) + (1 - \mu^G)u(d, d)$$

for $k \geq K$.

Proof. Since bad standing is absorbing, the value function for any bad standing score simply equals the flow payoff to bad-standing players. Hence, $V_k = \mu^G u(c, b) + (1 - \mu^G)u(d, d)$ for $k \geq K$.

For a good standing score k , the value function V_k is related to V_{k+1} , the value function at the next score $k + 1$, by

$$V_k = (1 - \gamma)(\mu^G u(a, a) + (1 - \mu^G)u(b, c)) + \gamma(1 - \nu(\varepsilon))V_k + \gamma\nu(\varepsilon)V_{k+1},$$

which gives

$$V_k = (1 - \alpha(\gamma, \varepsilon))(\mu^G u(a, a) + (1 - \mu^G)u(b, c)) + \alpha(\gamma, \varepsilon)V_{k+1}.$$

This, along with the fact that $V_K = \mu^G u(c, b) + (1 - \mu^G)u(d, d)$ implies the expression

for V_k when $k < K$. ■

OA.2 Proof of Supporting Results for Lemma 2

OA.2.1 Proof of Lemma 9

Lemma 9. *In a trigger strategy equilibrium with tolerance K and share of good-standing players μ^G , for all $k < K$,*

$$\mu_k = \beta(\gamma, \varepsilon)^k (1 - \beta(\gamma, \varepsilon)).$$

Proof. The inflow into score 0 is $1 - \gamma$, while the outflow from score 0 is $(1 - \gamma)\mu_0 + \gamma(n - 2)\varepsilon\mu^0$. Setting these equal gives

$$\mu_0 = \frac{1 - \gamma}{1 - \gamma + \gamma(n - 2)\varepsilon} = 1 - \beta(\gamma, \varepsilon).$$

Additionally, for every $0 < k < K$, both score k and score $k - 1$ are good standing scores. Thus, the inflow into score k is $\gamma(n - 2)\varepsilon\mu_{k-1}$, while the outflow from score k is $(1 - \gamma)\mu_k + \gamma(n - 2)\varepsilon\mu_k$. Setting these equal gives

$$\mu_k = \frac{\gamma(n - 2)\varepsilon}{1 - \gamma + \gamma(n - 2)\varepsilon} \mu_{k-1} = \beta(\gamma, \varepsilon) \mu_{k-1}.$$

Combining this with $\mu_0 = 1 - \beta(\gamma, \varepsilon)$ gives $\mu_k = \beta(\gamma, \varepsilon)^k (1 - \beta(\gamma, \varepsilon))$ for $0 \leq k \leq K - 1$. ■

OA.2.2 Proof of Lemma 10

Lemma 10. *In a trigger strategy equilibrium with tolerance K and share of good-standing players μ^G ,*

$$V_k = (1 - \beta(\gamma, \varepsilon)^{K-k}) (\mu^G u(a, a) + (1 - \mu^G) u(b, c)) + \beta(\gamma, \varepsilon)^{K-k} (\mu^G u(c, b) + (1 - \mu^G) u(d, d))$$

for $k < K$, and

$$V_k = \mu^G u(c, b) + (1 - \mu^G)u(d, d)$$

for $k \geq K$.

Proof. Since bad standing is absorbing, the value function for any bad standing score equals the flow payoff to bad-standing players, so $V_k = \mu^G u(c, b) + (1 - \mu^G)u(d, d)$ for $k \geq K$.

For a good standing score k , the value function V_k is related to V_{k+1} , the value function at the next score $k + 1$, by

$$V_k = (1 - \gamma)(\mu^G u(a, a) + (1 - \mu^G)u(b, c)) + \gamma(1 - (n - 2)\varepsilon)V_k + \gamma(n - 2)\varepsilon V_{k+1},$$

which gives

$$V_k = (1 - \beta(\gamma, \varepsilon))(\mu^G u(a, a) + (1 - \mu^G)u(b, c)) + \beta(\gamma, \varepsilon)V_{k+1}.$$

This, along with the fact that $V_K = \mu^G u(c, b) + (1 - \mu^G)u(d, d)$ implies the expression for V_k when $k < K$. ■

OA.3 Proof of Lemma 11

Lemma 11. *There is an equilibrium with share μ^G of players in G iff the following feasibility constraint*

$$\tau_b - \tau_a = \frac{1 - \gamma}{\gamma} h_n(\gamma, \varepsilon, \mu^G, \tau_a, \chi),$$

and incentive constraints

$$(a \text{ vs } b|a) : (1 - n\varepsilon)h_n(\gamma, \varepsilon, \mu^G, \tau_a, \chi)(\mu^G(u(a, a) - u(c, b)) + (1 - \mu^G)(u(b, c) - u(d, d))) \\ > u(b, a) - u(a, a),$$

$$(b \text{ vs } a|c) : (1 - n\varepsilon)h_n(\gamma, \varepsilon, \mu^G, \tau_a, \chi)(\mu^G(u(a, a) - u(c, b)) + (1 - \mu^G)(u(b, c) - u(d, d))) \\ < u(b, c) - u(a, c),$$

$$(a \text{ vs } x|a) : (1 - n\varepsilon)\frac{\gamma}{1 - \gamma}(\chi - \tau_a)\mu^G(\mu^G(u(a, a) - u(c, b)) + (1 - \mu^G)(u(b, c) - u(d, d))) \\ > \max_{x \neq a, b} u(x, a) - u(a, a),$$

$$(b \text{ vs } x|c) : (1 - n\varepsilon)\frac{\gamma}{1 - \gamma}(\chi - \tau_b)\mu^G(\mu^G(u(a, a) - u(c, b)) + (1 - \mu^G)(u(b, c) - u(d, d))) \\ > \max_{x \neq a, b} u(x, c) - u(b, c),$$

are satisfied.

Proof. The steady state condition for μ^G , the share of players in G , is

$$1 - \gamma = (1 - \gamma)\mu^G + \gamma(\mu^G((1 - (n - 1)\varepsilon)\tau_a \\ + \varepsilon\tau_b + (n - 2)\varepsilon\chi) + (1 - \mu^G)(\varepsilon\tau_a + (1 - (n - 1)\varepsilon)\tau_b + (n - 2)\varepsilon\chi))\mu^G.$$

This is equivalent to

$$\tau_b - \tau_a = \frac{1 - \gamma}{\gamma}h_n(\gamma, \varepsilon, \mu^G, \tau_a, \chi),$$

which gives the feasibility constraint.

By playing a , the expected continuation value of a player in G is $((1 - (n - 1)\varepsilon)(1 - \tau_a) + \varepsilon(1 - \tau_b) + (n - 2)\varepsilon(1 - \chi))V^G + ((1 - (n - 1)\varepsilon)\tau_a + \varepsilon\tau_b + (n - 2)\varepsilon\chi)V^B$, while the expected continuation value of a player in G from playing b is $(\varepsilon(1 - \tau_a) + (1 - (n - 1)\varepsilon)(1 - \tau_b) + (n - 2)\varepsilon(1 - \chi))V^G + (\varepsilon\tau_a + (1 - (n - 1)\varepsilon)\tau_b + (n - 2)\varepsilon\chi)V^B$. Thus, the incentive constraint for a good-standing player to play a rather than b against a is

$$(1 - n\varepsilon)\frac{\gamma}{1 - \gamma}(\tau_b - \tau_a)(V^G - V^B) > u(c, a) - u(a, a),$$

Since

$$V^G = \mu^G (\mu^G u(a, a) + (1 - \mu^G)u(b, c)) + (1 - \mu^G) (\mu^G u(c, b) + (1 - \mu^G)u(d, d))$$

and

$$V^B = \mu^G u(c, b) + (1 - \mu^G)u(d, d),$$

this inequality simplifies to

$$\begin{aligned} & (1 - n\varepsilon)h_n(\gamma, \varepsilon, \mu^G, \tau_a, \chi)(\mu^G(u(a, a) - u(c, b)) + (1 - \mu^G)(u(b, c) - u(d, d))) \\ & > u(b, a) - u(a, a), \end{aligned}$$

which gives the (a vs $b|a$) constraint. All the other incentive constraints can be similarly derived. ■

OA.4 Proof of Supporting Results for Theorem 5

Lemma 13. *For μ^G sufficiently close to 1, there is no sequence of bad-standing records $\{r_n\}_{n \in \mathbb{N}}$ such that $\lim_{n \rightarrow \infty} V_{r_n} = \bar{V}$.*

Proof. Suppose that there is a sequence of bad-standing records $\{r_n\}_{n \in \mathbb{N}}$ such that $\lim_{n \rightarrow \infty} V_{r_n} = \bar{V}$. Since $V_{r_n} = (1 - \gamma)\pi^B + \gamma p_r^c V_{r_n}^c + \gamma p_{r_n}^d V_{r_n}^d$ for all r_n , it follows that

$$V_{r_n} \leq \pi^B + \frac{\gamma}{1 - \gamma} p_r^c \max\{V_{r_n}^c - V_{r_n}, 0\} + \frac{\gamma}{1 - \gamma} p_{r_n}^d \max\{V_{r_n}^d - V_{r_n}, 0\}.$$

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 $\bar{V} = \lim_{n \rightarrow \infty} V_{r_n} \leq \pi^B$, so $V_r \leq \pi^B$ for all r .

This implies that $\pi^G \leq \pi^B$. However, since $\max_x u(x, a) > u(a, a)$, then for μ^G sufficiently close to 1, it would be strictly better for players to not follow the proposed trigger strategy, but rather, from entering the population on, always play their static best response to their opponent's action in every period, even when they are at good-

standing records. ■

Lemma 14. *If for every $\eta > 0$ there exist some γ and ε such that trigger strategies support share at least $1 - \eta$ of the population playing a , then $u(a, a) \geq u(c, b)$.*

Proof. Suppose towards a contradiction that $u(a, a) < u(c, b)$. Then there exists $\bar{\mu}^G < 1$ such that, whenever $\mu^G > \bar{\mu}^G$, the flow payoff to a bad-standing player strictly exceeds that of a good-standing player. In this case, since (a, a) is not a Nash equilibrium of the stage game, a good-standing player strictly prefers to play a best-response to a rather than playing a when matched with another good-standing player, which is a contradiction. Hence, trigger strategies cannot support more than share $\bar{\mu}^G$ playing a , regardless of γ and ε . ■

OA.5 Proof of Theorem 7

Fix parameters $\gamma, \varepsilon \in (0, 1)$ and suppose that all players play according to the PPR trigger strategy with transition probability parameters τ_a, τ_b , and χ , with $\tau_b \geq \tau_a$. Let $\mu^{t,G}$ denote the share of the population in good standing at time t . Define $p : (0, 1)^5 \rightarrow (0, 1)$ as $p(\varepsilon, \tau_a, \tau_b, \chi, \mu) = \mu^G((1 - (n - 1)\varepsilon)\tau_a + \varepsilon\tau_b + (n - 2)\varepsilon\chi) + (1 - \mu^G)(\varepsilon\tau_a + (1 - (n - 1)\varepsilon)\tau_b + (n - 2)\varepsilon\chi)$.

Given the value of $\mu^{t,G}$ at time t , $p(\varepsilon, \tau_a, \tau_b, \chi, \mu^{t,G})$ represents the probability that a good-standing player transitions to bad standing in the period. Hence, the corresponding value of $\mu^{t+1,G}$ is given by

$$\mu^{t+1,G} = 1 - \gamma + \gamma(1 - p(\varepsilon, \tau_a, \tau_b, \chi, \mu^{t,G}))\mu^{t,G}. \quad (6)$$

Lemma 19. *There is a unique steady-state μ_*^G .*

Proof. The steady state condition $1 - \gamma - \mu^G + \gamma(1 - p(\varepsilon, \tau_a, \tau_b, \chi, \mu^G))\mu^G = 0$ is quadratic in μ^G , and the left-hand side of the equation is $1 - \gamma > 0$ for $\mu^G = 0$ and $-p(\varepsilon, \tau_a, \tau_b, \chi, 1) < 0$ for $\mu^G = 1$. ■

We now complete the proof of Theorem 7. Note that $p(\varepsilon, \tau_a, \tau_b, \chi, \mu^{t,G})$ is decreasing in $\mu^{t,G}$. By Equation 6, $\mu^{t+1,G}$ is an increasing function of $\mu^{t,G}$. Let $\bar{\mu}^{t,G}$ be the time-path of the share of good-standing players when $\bar{\mu}^{0,G} = 1$, that is all players are in good standing at the initial period. Since $\mu^{t+1,G}$ is an increasing function of $\mu^{t,G}$, $\bar{\mu}^{t+1,G} \leq \bar{\mu}^{t,G}$ for all t . Thus, $\bar{\mu}^{t,G}$ converges and, since there is a unique steady state μ_*^G , it must be that $\lim_{t \rightarrow \infty} \bar{\mu}^{t,G} = \mu_*^G$. A similar argument shows that $\lim_{t \rightarrow \infty} \underline{\mu}^{t,G} = \mu_*^G$ where $\bar{\mu}^{t,G}$ is the time-path of the share of good-standing players when $\underline{\mu}^{0,G} = 0$. For any initial $\mu^{0,G}$, the monotonicity of $\mu^{t+1,G}$ in $\mu^{t,G}$ implies that $\underline{\mu}^{t,G} \leq \mu^{t,G} \leq \bar{\mu}^{t,G}$ for all t , so $\lim_{t \rightarrow \infty} \mu^{t,G} = \mu_*^G$.