

Steady-State Equilibria in Anonymous Repeated Games*

Daniel Clark, Drew Fudenberg, and Alexander Wolitzky

March 10, 2020

Abstract

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

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

*This paper replaces our earlier papers “Steady-State Equilibria in Anonymous Repeated Games, I: Trigger Strategies in General Stage Games,” “Steady-State Equilibria in Anonymous Repeated Games, II: Coordination-Proof Strategies in the Prisoner’s Dilemma,” and “Robust Cooperation with First-Order Information.” We thank Nageeb Ali, V Bhaskar, Glenn Ellison, Sander Heinsalu, Yuval Heller, Takuo Sugaya, Satoru Takahashi, and Caroline Thomas for helpful comments and conversations, and NSF grants SES 1643517 and 1555071 and Sloan Foundation grant 2017-9633 for financial support.

1 Introduction

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

To study these situations, we introduce a new class of repeated games with anonymous random matching. There is a unit mass of players with geometrically distributed lifespans and a constant inflow of new players that keeps the total population constant. 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 investigate how steady-state cooperation can be supported under various sorts of “record systems,” which provide each player with some information about the past behavior or standing of their current opponent. We ask what types of record systems allow the community to support good outcomes, both for general stage games and for the leading special case where the stage game is the prisoner’s dilemma (PD).

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 statistically identify the vector of their opponents’ actions, equilibria that support 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

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.

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.

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 is “finite partitional,” which means that for any given record, there is a finite partition of the opponent’s record space such that the update function does not vary with the opponent’s record within each 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.

Our main results characterize the prospects for cooperation under different types of record systems. We consider only strict equilibria; this captures a simple form of robustness, and in particular rules out “belief-free” equilibria and related constructions. For most of our results, we also require equilibria to be *coordination-proof*: a pair of matched players never play a Pareto-dominated equilibrium in the “augmented” game induced by their current records and their expected continuation payoffs. Restricting attention to coordination-proof strategies rules out equilibria built on within-match

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).

miscoordination. Finally, we focus on the double limit where players' expected lifespans are long and there is little noise in the record system. Taking this limit allows a clean analysis and gives society its best shot at supporting cooperative outcomes; of course, if players are myopic or records are extremely noisy, only static Nash equilibrium outcomes can arise.

We begin by analyzing *second-order records*, which allow record updates to depend on a player's own record and action as well as their partner's action, but not their partner's record. Thus, a player's second-order record updates based solely on the outcomes of their own interactions, without requiring information about the outcomes of their partner's past interactions, so second-order records are more informationally robust than more general interdependent records. We show that second-order records are rich enough to support the play of any action that Pareto-dominates the pure-strategy minmax payoff (in the long-lifespans, low-noise double limit).³ To prove this, we consider strategies that assign players to good or bad standing based on their records, and specify that good-standing players take the target action when matched with each other, but players take a minmaxing action whenever at least one player in the match has bad standing. With these strategies, second-order records can identify the standing of a good-standing player's partner from their action, which allows the threat of switching a good-standing player to bad standing to incentivize the prescribed behavior among good-standing players.⁴

We then turn to *first-order records*, which require that a player's record is updated based only on their own record and action. These records capture situations where the available information about one's current partner depends on their past actions, but not on their past partners' actions. For example, in online marketplaces such as Ebay or AirBnB, users typically rate their current partner's behavior in the absence of any

³The equilibrium we construct to prove this "minmax-threat" folk theorem is strict but may fail to be coordination-proof. We also provide a "Nash-threat" folk theorem based on equilibria that are both strict and coordination-proof.

⁴Likewise, bad-standing players can be incentivized by the promise of an eventual return to good standing.

information about the current partner’s past partners’ behavior. Future users then observe summary statistics of these ratings, which is a form of first-order information.

First-order record systems cannot support as many actions as second-order systems can, because first-order records cannot distinguish “justified” deviations from the target action from “unjustified” ones. For example, in the PD, if a player is penalized for defecting against opponents who cooperate, they must also be penalized for defecting against opponents who defect. This makes it much harder to support a steady state with a high share of cooperation.

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

Our positive results for second-order records and for first-order records with unprofitable punishments beg the question of when an action without an unprofitable punishment can be supported with first-order records. The remainder of our analysis answers this for the leading example of cooperation in the PD: We characterize the set of payoff parameters in the PD for which there exist strict, coordination-proof equilibria in which the share of cooperation converges to 1 in the long-lifespans, low-noise

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

double limit.⁶ We then show that for a complementary set of payoff parameters, the only strict, coordination-proof equilibrium is *Always Defect*.

To understand our characterization, recall the standard normalization of the PD payoff matrix, where $g, l > 0$ and $g < l + 1$, so (C, C) maximizes the sum of payoffs.

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

Figure 1: The Prisoner’s Dilemma

Here g measures the gain to defection when one’s opponent cooperates, while l measures the gain to defection when one’s opponent defects. The comparison of g and l thus reflects the degree of complementarity in the PD, which we will see is an important factor in determining whether equilibrium cooperation is possible. Intuitively, with first-order records a player’s gain in continuation payoff from playing C rather than D is independent of their opponent’s action. Hence, a player can be strictly incentivized to play C against opponents who cooperate while playing D against opponents who defect only if the short-term cost of playing C rather than D is greater when the opponent defects. For this reason, cooperation can be supported in a strict equilibrium only if $g < l$.

We say that the game is *submodular* when $g \geq l$, *strictly supermodular* when $g < l$, and *strongly supermodular* when $g + g^2 < l$. We also say it is *mild* when $g < 1$ and *severe* when $g \geq 1$. We find that full (limit) cooperation is possible in strict, coordination-proof equilibrium if the game is either strongly supermodular or both strictly supermodular and mild, and that otherwise no cooperation is possible.

When the game is strictly supermodular and mild, cooperation can be supported by strategies where good-standing players cooperate with each other and defect against bad-standing players, bad-standing players always defect, and players cycle between good and bad standing depending on the number of times they have defected. A rough

⁶The order of limits turns out to matter for this result: here we must first take expected lifespans to infinity, and then take noise to zero.

intuition for why these “cyclic” strategies form an equilibrium when $g < 1$ but not when $g \geq 1$ is that, since the immediate gain from playing D rather than C against C is g , and the future loss from switching the outcome of a match from (C, C) to (D, D) is 1, incentives require that each play of D leads to at least g “switches” from (C, C) to (D, D) . Intuitively, when $g < 1$ (the game is mild), these switches “peter out,” allowing a steady state with positive cooperation, but when $g \geq 1$ (the game is severe), the switches “snowball,” which precludes positive steady-state cooperation under any strategies where only (C, C) and (D, D) are played on path.⁷

When the game is severe, since no strategy where only (C, C) and (D, D) are played on path can support cooperation, any equilibrium with positive cooperation must involve some on-path plays of (C, D) : that is, players with certain records must be willing to cooperate even when their opponent is expected to defect. We call players with such records *supercooperators*, and we show that equilibria with supercooperators exist (in the double limit) if and only if the game is strongly supermodular: $g + g^2 < l$. Intuitively, higher values of g make it harder to support cooperation because deterring defection requires harsher punishments when g is larger, and these punishments occur with positive frequency on path. In contrast, higher values of l help support cooperation by preventing undesired cooperation against defectors.⁸ The threshold $g + g^2 = l$ for the emergence of supercooperation comes from combining such incentive effects with the equations for steady-state population shares.

1.1 Related Work

Rosenthal (1979) and Rosenthal and Landau (1979) introduced repeated games with random matching. Rosenthal (1979) considered finite populations in the special case of first-order information where players observe only their current opponent’s most recent

⁷This argument imagines the dynamic consequences of a single play of D . Our formal arguments instead directly analyze the equations that define a steady-state equilibrium.

⁸This force is similar to that in Kandori (1992), but Kandori’s construction requires deterring all supercooperation, while our construction requires a positive level of supercooperation.

action, and proved that Markovian equilibria exist.⁹

Kandori (1992) and Ellison (1994) show contagion equilibria can support cooperation in the finite-population PD 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.¹⁰

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 a continuum-population model in the absence of noise.¹¹ Takahashi (2010) shows how cooperation can be supported in the PD when players observe their partner’s entire past history of actions—all first-order information—but no higher-order information. For the submodular PD, his construction relies on belief-free mixed strategies, which our restriction to strict equilibria rules out. For the supermodular PD, he constructs a strict equilibrium in the absence of noise, based on intricate conditioning on calendar time.¹² Equilibria of this form cannot exist in our model, since we allow noise and players do not share a common notion of calendar time.

Heller and Mohlin (2018) suppose players observe a finite sample of their current partner’s past actions, a particular form of first-order information. Players live forever, 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. A small fraction of players are assumed to be commitment types, so that a partner’s past actions are a noisy signal of their type and thus of their likely current action. This allows the

⁹Rosenthal and Landau (1979) studied 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.

¹⁰Ellison extended Kandori’s results to allow both for arbitrary payoff parameters and for a small amount of noise. Deb, Sugaya, and Wolitzky (2018) generalize these results to prove the folk theorem for finite-population repeated games with anonymous random matching. Their construction relies heavily on a finite population, common calendar time, and non-strict incentives.

¹¹Sugden (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). Related notions of steady-state equilibria arise in papers on industry dynamics (Adlakha, Johari, and Weintraub, 2015; Hopenhayn, 1992; Jovanovic and Rosenthal, 1988)

¹²Specifically, his construction relies on “threading,” introduced by Ellison (1994).

construction of an efficient mixed-strategy equilibrium in the supermodular PD.^{13,14}

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

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. However, they study non-equilibrium learning as opposed to equilibrium cooperation.

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 $u : A \times A \rightarrow \mathbb{R}$.

This section presents our model of records, states, and steady-state equilibria.

¹³Heller and Mohlin also consider alternative information structures where players observe, for example, a finite sample of their partners’ past action *profiles*. Dilmé (2016) also assumes commitment types, and constructs an efficient belief-free equilibrium for the case where $g = l$.

¹⁴Bhaskar and Thomas (2018) study first-order information in a sequential-move game with one-sided incentive constraints. The distinction between submodular and supermodular games does not arise due to the sequential nature of the game.

2.1 Record Systems

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

Definition 1. A *record system* \mathcal{R} is a 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**).

Note that the update rule is allowed to be stochastic. This can capture errors in recording, as well as imperfect implementation of the intended action. 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 record update does not depend directly on their opponent’s record.

Definition 2. A record system is **second-order** if the update rule can depend on both players’ actions but only the player’s own record. Here $\mathcal{D} = R \times A^2$.

With a second-order record system, a player’s record can be computed based only on their own history of stage-game outcomes.¹⁵

¹⁵If contrary to our assumptions the record system is “centralized” in the sense of having access to

Definition 3. A record system is **first-order** if the update rule can depend only on the player’s own action and record. Here $\mathcal{D} = R \times A$.

Nowak and Sigmund (1998), Takahashi (2010), Bhaskar and Thomas (2018), and Heller and Mohlin (2018) also consider first-order records.¹⁶ We consider second-order records in Section 3 and first-order records in Sections 4 and 5.

2.2 Strategies, States, and Steady States

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

Thus, all strategies that condition only on payoff-relevant variables are *pairwise-public*, meaning that they condition only on information that is public knowledge between the two partners, namely their records.¹⁷ We restrict attention to such strategies. We write a pairwise-public pure strategy as functions $s : R \times R \rightarrow A$, with the convention that the first coordinate is the player’s own record and the second coordinate is the partner’s record, and similarly write a pairwise-public mixed strategy as a function $\sigma : R \times R \rightarrow \Delta(A)$. We also assume that all players use the same strategy. Note that every strict equilibrium in a symmetric, continuum-population model is pairwise-public and symmetric, so these restrictions are without loss for strict equilibria.

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 is μ , the distribution of matches is given by $\mu \times \mu$.

all players’ records, then sufficiently detailed second-order records could be used to compute and track the “status levels” implied by more general interdependent records. However, this is not true for the decentralized records we consider (where the update domain is a subset of $R^2 \times A^2$). Kocherlakota (1998) makes a similar point in the context of macroeconomic matching models.

¹⁶To interpret noisy first-order records as resulting from implementation errors, the outcome of the game must have a product structure in the sense of Fudenberg, Levine, and Maskin (1994), so that a player’s record update does not depend on the opponent’s action.

¹⁷Rosenthal (1979) called such strategies “Markovian.”

That is, for each $(r, r') \in R^2$ with $r \neq r'$, the fraction of matches between players with record r and r' is $2\mu_r\mu_{r'}$, while the fraction of matches between two players with record r is μ_r^2 .

Given a record system \mathcal{R} and a pairwise-public strategy σ used by all players, denote the distribution over next-period records of a player with record r who meets a player with record r' by

$$\phi_{r,r'}(\sigma) = \sum_a \sum_{a'} \sigma(r, r')[a] \sigma(r', r)[a'] \rho(r, r', a, a') \in \Delta(R).^{18}$$

The *state update map* $f_\sigma : \Delta(R) \rightarrow \Delta(R)$ is then given by

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

(Here the $(1 - \gamma)$ term in the formula for $f_\sigma(\mu)[0]$ reflects the per-period inflow of new players, who enter the game with record 0.) A *steady state* under σ is a state μ such that $f_\sigma(\mu) = \mu$.

Theorem 1.

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

¹⁸This equation applies for general interdependent records. To apply it to more restrictive 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.

The proof is in Appendix A.1; all other omitted proofs can also be found in either the Appendix (A) or the Online Appendix (OA). Intuitively, the combination of the finite domain of the record-update function (due to finiteness of the stage game and, for interdependent record systems, the finite-partition property) and geometrically distributed lifetimes imply that almost all player’s records lie in a finite subset of the record set. This lets us find a set \bar{M} that contains all feasible distributions over records and resembles a finite-dimensional set—in particular, \bar{M} is compact in the sup norm. We then show that f maps \bar{M} to itself and is continuous in the sup norm so, since \bar{M} is also convex, we can appeal to a fixed point theorem.¹⁹ When instead the record-update function does not have a 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 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.²⁰ Intuitively, this multiplicity corresponds to different initial conditions at time $t = -\infty$.

2.3 Steady-State Equilibria

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

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

¹⁹Fudenberg and He (2018) use a similar approach to establish steady-state existence in a random matching model with geometric lifetimes and countable records. In their model players do not observe each other’s records, so the finite-partition property (which that paper does not discuss) is automatically satisfied.

²⁰For instance, suppose that $R = \{0, 1, 2\}$, the action set is singleton, and newborn players have record 0. When matched with a player with record 0 or 1, the record of a player with record 0 or 1 increases by 1 with probability ε and remains constant with probability $1 - \varepsilon$, but it increases by 1 with probability 1 when the player is matched with a player with record 2. When a player’s record reaches 2, it remains 2 for the remainder of their lifetime. Depending on the parameters γ and ε , there can be between one and three steady states in this example.

where we have extended the payoff function u to mixed strategies in the usual way. Next, denote the probability that a player with record r today has record r' t periods from now (assuming they are still alive) 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.

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

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

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

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 static equilibrium α^* , and let σ recommend α^* at every record

pair (r, r') . Then (σ, μ) is an equilibrium for any steady state μ . ■

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

Corollary 2. *Under any 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 static equilibrium.

2.4 Coordination-Proofness

We will frequently impose the equilibrium refinement of *coordination-proofness*. This captures the idea that equilibria that rely on “miscoordination” within a match will break down if matched partners manage to coordinate successfully.

For a fixed equilibrium (σ, μ) , denote the expected continuation payoff of a player with record r who plays action a against an opponent with record r' who plays a' by $V_{r,r'}^{a,a'} := \sum_{r''} \rho(r, r', a, a') [r''] V_{r''}$. The *augmented payoff function* $\hat{u} : R \times R \times A \times A \rightarrow \mathbb{R}$ is defined by $\hat{u}_{r,r'}(a, a') := (1 - \gamma)u(a, a') + \gamma V_{r,r'}^{a,a'}$. The *augmented game* between players with records r and r' is the static game with action set $A \times A$ and payoff functions $\hat{u}_{r,r'}$ and $\hat{u}_{r',r}$.

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

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

The logic is that, if $(\sigma(r, r'), \sigma(r', r))$ were Pareto-dominated by some augmented-game Nash equilibrium (α, α') , players with records r and r' would benefit from reaching a self-enforcing agreement to play (α, α') when matched together, breaking the equilibrium.²¹

A sufficient condition for the existence of a coordination-proof equilibrium is that the stage game admits a symmetric Nash equilibrium that is not Pareto-dominated by another (possibly asymmetric) Nash equilibrium. This condition is satisfied in many games, including the PD.

Corollary 3. *Under any first or second-order record system, a coordination-proof equilibrium exists if the stage game has a symmetric Nash equilibrium that is not Pareto-dominated by another Nash equilibrium. The same is true for any finite-partitional interdependent record system.*

Proof. Fix such a symmetric static equilibrium α^* , and let σ recommend α^* at every record pair (r, r') . Then (σ, μ) is an equilibrium for any steady state μ . Moreover, note that $\hat{u}_{r,r'}(a, a') = (1 - \delta)u(a, a') + \delta u(\alpha^*, \alpha^*)$, for any r, r', a, a' . Thus, (α, α') is a (possibly mixed) augmented-game Nash equilibrium if and only if it is a Nash equilibrium of the stage game. Since (α^*, α^*) is not Pareto-dominated by another static equilibrium, there is no augmented-game Nash equilibrium (α, α') satisfying $(u(\alpha, \alpha'), u(\alpha', \alpha)) > (u(\alpha^*, \alpha^*), u(\alpha^*, \alpha^*))$, and hence there is no augmented-game Nash equilibrium (α, α') satisfying $(\hat{u}_{r,r'}(\alpha, \alpha'), \hat{u}_{r',r}(\alpha', \alpha)) > (\hat{u}_{r,r'}(\alpha^*, \alpha^*), \hat{u}_{r',r}(\alpha^*, \alpha^*))$ for any r, r' . That is, (σ, μ) is coordination-proof. ■

2.5 Canonical Records

Our positive results require that the amount of noise in the record system is sufficiently small. To formalize this, we introduce the notion of a canonical record system, which

²¹Coordination-proofness is somewhat reminiscent of renegotiation-proofness in fixed-partner repeated games as studied by Farrell and Maskin (1989) and others, but it is simpler since each pair of partners plays a single one-shot game.

is one that noisily records every action (in the case of first-order records) or every stage-game outcome (for second-order records).

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

Definition 5.

1. A first-order record system is **canonical** if the record set R is canonical and there exists a noise matrix ε such that, for every record $r = \prod_{\tau=1}^t a_{\tau}$ and action a , we have

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

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

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

In general, the set of equilibria depends on both the amount of noise in the system and the players' expected lifetimes. We focus on the case where there is vanishingly

little noise and players live a long time: that is, on the double limit $(\gamma, \varepsilon) \rightarrow (1, 0)$, where ε is the noise matrix in the canonical record-system, and $\varepsilon \rightarrow 0$ means that every entry in the matrix ε converges to 0.

We consider two ways of passing to this limit.

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

Some of our results apply to the general double limit, while others hold only for the iterated limit. This iterated limit matches the order of limits in related papers on almost-perfect private monitoring (e.g., Ellison (1994), Hörner and Olszewski (2006)). We will also ask when an action can be limit-supported by strict and/or coordination-proof equilibria. Here the same definition applies with the appropriate equilibrium refinement.

3 Second-Order Records: Folk Theorems

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

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

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

- (i) If $u(a, a) > \underline{u}$, then a can be limit-supported by strict equilibria.

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

Theorem 2(i) is a minmax-threat folk theorem. The construction relies on “cyclic” strategies of the following form: Let $b \in \arg \min_{a'} \max_a u(a, a')$ be a minmax action. Players begin in good standing. A player in good standing plays a when matched with a fellow good-standing player and plays b against bad-standing players, while a player in bad standing always plays b . When a good-standing player’s outcome is recorded as something other than (a, a) or (b, b) , the player enters bad standing. A player remains in bad standing until they accumulate M (b, b) profiles for some fixed $M \in \mathbb{N}$, at which point they return to good standing. We show that, when $(\gamma, \varepsilon) \approx (1, 0)$, M can be chosen to be high enough so that the punishment of M periods of (b, b) is severe enough to deter deviations from the prescribed strategy, but also low enough that the steady-state share of players in good standing is high.

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

Theorem 2(ii) presents a condition under which an action a can be iterated-limit-supported by strict, coordination-proof equilibria.²² It gives a Nash-threat folk theorem, where the “threat point” equilibrium (b, b) is required to be strict, symmetric, and not Pareto-dominated by another static equilibrium. For example, in the prisoner’s dilemma, taking $a = C$ and $b = D$ implies that *Cooperate* is iterated-limit-supported by strict, coordination-proof equilibria.

The proof of this part relies on “tolerant grim trigger strategies.” For each player,

²²Thus, the second part of the theorem is stronger than the first in that it imposes coordination-proofness, but weaker in that it covers a smaller set of actions a and establishes only that these actions are iterated-limit-supported, rather than limit-supported.

define their *score* $k \in \mathbb{N}$ as the number of times their outcome was recorded as anything other than (a, a) or (b, b) . 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 against fellow good-standing players and play b against bad-standing players, while bad-standing players always play b .²³ 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 independently 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.

Intuitively, the required properties can be satisfied by choosing K to be slightly higher than the highest score a newborn player expects to attain in their lifetime. By the law of large numbers, a player who follows the equilibrium strategy is unlikely to reach such a score K , so the steady-state share of players in good standing is close to 1. At the same time, the probability of reaching score K during one's lifetime increases quickly enough in one's current score to deter deviations.

Okuno-Fujiwara and Postlewaite (1995)'s Theorem 1 shows that with a form of interdependent records (what they call "status"), any actions that Pareto-dominate the pure-strategy minmax payoffs can be supported in equilibrium without noise. Their proof uses standard intolerant grim trigger strategies, so it is not robust to noise. Theorem 2 shows that their theorem's conclusion does not require interdependent records, and that by tuning the tolerance level it also extends to settings with overlapping generations and noise.

²³Tolerant grim strategies were introduced by Fudenberg, Rand, and Dreber (2012) in the context of repeated prisoner's dilemma games with fixed partners. Some experimental subjects seem to use such strategies.

4 First-Order Records: Unprofitable Punishments

Now we turn to first-order record systems, where the updating of a player’s record depends only on their own play. Such records cannot support as many actions as second-order records can, and the folk theorem fails for strict equilibrium. The key obstacle is that first-order records cannot distinguish “justified” deviations from the target action profile from “unjustified” ones. For example, in the PD, if players are penalized for playing *Defect* against *Cooperate* (an off-path, opportunistic deviation), they must also be penalized for playing *Defect* against *Defect* (a justified punishment that must sometimes occur on-path if defection is to be deterred). This can preclude the existence of equilibria with a high share of players in good standing.

In this section, we note that this obstacle does not arise when the target action profile (a, a) has the property that there exists a punishing action b and a strict best response c to b such that $u(a, a) > u(c, b)$ (so that facing b is indeed a punishment), $u(a, a) > u(b, a)$ (so that deviating from a to b is unprofitable for a player whose opponent plays a), and $u(b, c) > u(a, c)$ (so a player prefers to carry out the punishment b rather than playing the target action, when the opponent best-responds to b). We say that in this case b is an *unprofitable punishment* for a . Intuitively, when an unprofitable punishment b exists for action a , the threat of switching to b can motivate one’s opponents to play a , but a player is not tempted to unilaterally deviate to b against opponents who play a . This enables first-order records to support the play of (a, a) by penalizing players only for taking actions other than a or b . In contrast, when the only punishing action b is a tempting deviation against a (as in the PD, where the punishing action D is always tempting), players must be penalized for playing b , and the record system’s inability to distinguish justified and unjustified plays of b becomes a real obstacle.

To establish this result, we assume that all possible recording errors are equally likely: $\varepsilon_{a,a'} = \varepsilon_{\tilde{a},\tilde{a}'}$ for all $a, a', \tilde{a}, \tilde{a}'$.²⁴ This uniform noise assumption isolates the payoff

²⁴This can be relaxed to assuming that the $\varepsilon_{a,a'}$ ’s lie in a small open ball around a uniform profile

consequences of different actions; without it, the required definition of an unprofitable punishment would be complicated by the need to account for differences in the $\varepsilon_{a,a'}$'s.

Theorem 3. *Fix an action a . With canonical first-order records and uniform noise:*

- (i) *If there exists an unprofitable punishment b for a and there is a strict and symmetric static equilibrium (d, d) that is not Pareto-dominated by another static equilibrium, then a can be limit-supported by strict equilibria.*
- (ii) *If there exists an action b such that (b, b) is a strict static equilibrium that is not Pareto-dominated by another static equilibrium and $u(a, a) > \max\{u(b, a), u(b, b)\}$, then a can be limit-supported by strict equilibria.*

The proof, which is in OA.2, is similar to the proof of Theorem 2(ii), except that now a player's score increases whenever they take an action other than a or b .²⁵

Note that the condition in Theorem 3(ii) applies when $b = c = d$ in the definition of an unprofitable punishment, in which case (a, a) can be supported by Nash reversion to (b, b) . For example, suppose the stage game is a PD 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 $u(D, D)$. Here both (E, E) and (D, D) are static equilibria, but E is not a profitable deviation against C , unlike D . Thus, Theorem 3 implies that cooperation can be limit-supported by Nash reversion to (E, E) .²⁶

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 of errors.

²⁵There are two further differences. First, the equilibria used to prove Theorem 3 may not be coordination-proof, because the action profile (b, c) prescribed in a match between a good-standing player and a bad-standing player may be Pareto-dominated by another Nash equilibrium in the augmented game between such players. Second, the proof of Theorem 3 handles the general $(\gamma, \varepsilon) \rightarrow (1, 0)$ limit rather than the iterated limit; this strengthened conclusion is an incidental benefit of the uniform noise assumption.

²⁶Technically, for E to satisfy the definition of an unprofitable punishment, we need to specify that $u(C, E)$ is strictly less than $u(E, E)$. If instead $u(C, E) = u(E, E)$, the same proof applies, but the constructed equilibria are no longer strict.

ment from an unwarranted deviation is one factor that has led Boyd et al. (2003), Gintis et al. (2003), and Bowles and Gintis (2011) (among others) to argue that the enforcement of human cooperation cannot be explained without appealing to social preferences. 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 cooperation in the PD is always limit-supported with second-order records, but (as we will see) is limit-supported with first-order records only for certain parameters, supports the argument that the inability to distinguish justified and unjustified plays of *Defect* is a serious obstacle to cooperation in the PD. However, this obstacle evaporates when a simple exit option is added to the game, consistent with the position of Baumard and Guala.

Another important special case of unprofitable punishments 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. 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 First-Order Records: Cooperation in the PD

For games in which unprofitable punishments do not exist, characterizing which actions can be limit-supported with first-order records is much more challenging. In this section, we resolve this question for the leading case of cooperation in the PD: we characterize the set of payoff parameters g and l for which cooperation can be limit-supported or iterated-limit-supported by strict, coordination-proof strategies in the PD. (Throughout, g and l refer to the payoff parameters in Figure 1.)

We first introduce some preliminary concepts in Sections 5.1 and 5.2. We then

present our main characterization result in Section 5.3. The characterization result involves a couple technical subtleties, which Section 5.4 shows can be resolved by introducing a type of randomization device.

5.1 Defectors, Supercooperators, Preciprocators

We begin with some terminology for different types of records.

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

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

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

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

Lemma 1. *Fix any first-order record system. In any strict equilibrium:*

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

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

1. Suppose that $g \geq l$. When two players with the same record r meet, by symmetry (an implication of strictness) they play either (C, C) or (D, D) . In the former case, C is the strict best response to C . Since the current-period gain from playing D instead of C is weakly smaller when the opponent plays D , this means C is also the strict best response to D , so record r is a supercooperator. In the latter case, D is the strict best response to D , and hence is also the strict best response to C , so record r is a defector.

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

Theorem 4. *Fix any first-order record system. If $g \geq l$, the unique strict equilibrium is Always Defect: $s(r, r') = D$ for all $r, r' \in R$.²⁷*

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

Theorem 4 confirms that strictly individually rational actions are not always limit-supportable by strict equilibria with first-order records, in contrast to the situation with second-order records. The rest of this section analyzes for what parameters g and l cooperation is limit-supportable by strict, coordination-proof equilibria with first-order records. Such parameters must of course satisfy $g < l$: that is, the PD must be strictly supermodular.

²⁷The conclusion of Theorem 4 extends to all (possibly non-strict) pure-strategy equilibria whenever $g > l$. Takahashi (2010) and Heller and Mohlin (2018) obtain the same conclusion in related models.

5.2 Coordination-Proofness in the PD

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

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

Proof. By definition, preciprocators play C against opponents who play C and play D against those who play D . Hence, the augmented game between any two preciprocators is a coordination game, with Nash equilibria (C, C) and (D, D) . Since playing D always gives a short-run gain, the fact that preciprocators play C against C implies that cooperating leads to higher continuation payoffs. Therefore, the (C, C) equilibrium yields both higher stage-game payoffs and higher continuation payoffs than the (D, D) equilibrium. That is, the (D, D) augmented-game equilibrium is Pareto-dominated by the (C, C) augmented-game equilibrium, so any pair of matched preciprocators must play (C, C) rather than (D, D) . ■

Coordination-proofness thus implies that every preciprocator plays C when matched with another preciprocator or a supercooperator, and plays D when matched with a defector. In particular, all preciprocators play C against the same set of opposing records. Hence, a strict, coordination proof equilibrium is completely characterized by a description of which records are preciprocators, which are supercooperators, and which are defectors. Denote the total population share in these classes by μ^P, μ^S , and μ^D respectively. We will use the term *cooperator* for all players who are either preciprocators or supercooperators (i.e., anyone who is not a defector), and we denote the population share of cooperators by $\mu^C = \mu^P + \mu^S = 1 - \mu^D$.

5.3 Cooperation and Limit Efficiency

We now present necessary and sufficient conditions for cooperation to be iterated-limit-supportable in strict, coordination-proof equilibria with first-order records. Our sufficient conditions require canonical records with $\varepsilon \rightarrow 0$; however, it will become clear

that the same sufficient conditions apply under coarser record systems that track only the number of times a player has defected (again with vanishing noise). Our necessary conditions apply for any “noisy” first-order record system.

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

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

Here $q_C(r)$ represents the distribution over records after “a recording of C is fed into the record system,” $q_D(r)$ represents the distribution over records after “a recording of D is fed into the record system,” and the ε ’s represent noise. The key feature of this definition is that perfect recording of actions is ruled out by the assumption that $\varepsilon_C(r) > 0$.

Recall that the prisoner’s dilemma is *mild* if $g < 1$ and *severe* otherwise, and that the game is *strongly supermodular* if $l > g + g^2$. Define the function $f : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ by $f(g) = \min \left\{ \frac{1+g}{|\ln(1+g)-1|}, 22.0625 - 3.63636g \right\}$.

We obtain the following characterization.

Theorem 5.

- (i) *With any noisy first-order record system, if $g \geq 1$ and $l \leq g + g^2$ (i.e., the prisoner’s dilemma is severe and not strongly supermodular), the only strict, coordination-proof equilibrium is Always Defect.*
- (ii) *With canonical first-order records, if either $g < 1$ or $l > \max\{g + g^2, f(g)\}$, cooperation can be iterated-limit supported by strict, coordination-proof equilibria.*

The proof uses the slack provided by the function f to rule out parameters for which we do not know how to construct sequences of iterated-limit-efficient strategies given the discreteness of players’ records. Note though that $f(g) < g + g^2$ whenever $g > 2.9$,




Figure 2: Iterated-limit efficiency obtains in the green region. In the blue region, limit efficiency obtains with PPR, but our results do not establish whether it obtains without PPR. In the red region, the only strict, coordination-proof equilibrium is *Always Defect*.

in which case f drops out of the theorem statement. Moreover, Section 5.4 proves the following result, which shows that if we introduce *personal public randomization* (PPR), a type of randomization device, the function f is not needed. Personal public randomization also allows the share of cooperators to converge to 1 in the general $(\gamma, \varepsilon) \rightarrow (1, 0)$ limit.

Theorem 6. *When personal public randomization is available,*

- (i) *With any noisy first-order record system, if $g \geq 1$ and $l \leq g + g^2$, the only strict, coordination-proof equilibrium is Always Defect.*
- (ii) *With canonical first-order records, if either $g < 1$ or $l > g + g^2$ (i.e., the prisoner's dilemma is either mild or strongly supermodular), cooperation can be limit-supported by strict, coordination-proof equilibria.*

Figure 2 displays the conclusions of Theorems 5 and 6. Note that as g increases from just below 1 to just above 1, the critical value of l above which cooperation is possible jumps from 1 to at least 2.

We now discuss the intuition for the necessary and sufficient conditions in Theorems 5 and 6. The proofs are contained in A.3 and OA.5, respectively.

5.3.1 Necessary Conditions

Section 1 discussed the intuition for why small values of g and large values of l make supporting cooperation easier. The specific necessary condition $g < 1$ or $l > g + g^2$ comes from combining two inequalities: $\mu^S < 1/(1 + g)$ and $\mu^P + \mu^S(l - g) > g$. The next lemma shows that combining these inequalities delivers the necessary condition $g < 1$ or $l > g + g^2$. After the lemma's short proof, we explain why the inequalities hold.

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

Proof. Suppose that $\mu^S < 1/(1 + g)$. Then the highest value of $\mu^P + \mu^S(l - g)$ is bounded above by either 1, which corresponds to $\mu^P = 1$ and $\mu^S = 0$, or $l/(1 + g)$, which corresponds to $\mu^P = g/(1 + g)$ and $\mu^S = 1/(1 + g)$. Hence, if $\mu^P + \mu^S(l - g) > g \geq 1$, then $l/(1 + g) > g$, which requires $l > g + g^2$. ■

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

The inequality $\mu^P + \mu^S(l - g) > g$ is shown by Lemma 10 in A.3.2.1. Here is a heuristic derivation in the special case $\mu^S = 0$, where the inequality simplifies to $\mu^P > g$. Note that this implies that $\mu^S > 0$ in any strict, coordination-proof equilibrium with positive cooperation when $g \geq 1$.

Fix a strict, coordination-proof equilibrium with $\mu^S = 0$. In any such equilibrium,

the only action profiles played on path are (C, C) and (D, D) . Hence, for any record r ,

$$\begin{aligned} V_r &= (1 - \gamma) \sum_{t=0}^{\infty} \gamma^t \mathbb{E} [\text{payoff in } t \text{ periods} \mid \text{record } r \text{ today}] \\ &= (1 - \gamma) \sum_{t=0}^{\infty} \gamma^t \Pr(\text{outcome } (C, C) \text{ in } t \text{ periods} \mid \text{record } r \text{ today}). \end{aligned}$$

Now assume that record r is a reciprocator. When $\varepsilon = 0$, the incentive constraint that a record r player takes C against an opponent who takes C is

$$\begin{aligned} (1 - \gamma)(1) + \gamma V_{r,C} &> (1 - \gamma)(1 + g) + \gamma V_{r,D}, \text{ or} \\ V_{r,C} - V_{r,D} &> \frac{1 - \gamma}{\gamma} g, \end{aligned}$$

where $V_{r,C}$ and $V_{r,D}$ denote the expected continuation payoffs of record r players who are recorded as playing C and D , respectively. (When $\varepsilon \neq 0$, this constraint is tighter.)

In words, this inequality says that a reciprocator who plays D instead of C must cause at least g/γ future outcomes to switch from (C, C) to (D, D) .

We now show that combining this incentive constraint with the steady-state condition $V_0 = (\mu^P)^2$ (which holds in any coordination-proof equilibrium with $\mu^S = 0$) implies that $\mu^P > g$. For simplicity, assume there is a “best” record $r^* = \arg \max_r V_r$, and that a player with record r^* is a reciprocator. Then $V_{r^*} \geq V_{r^*,C}$, so we have $V_{r^*} - V_{r^*,D} > (1 - \gamma)g/\gamma$. At the same time, when $\varepsilon = 0$ the Bellman equation for V_{r^*} is

$$\begin{aligned} V_{r^*} &= \mu^P (1 - \gamma + \gamma V_{r^*,C}) + (1 - \mu^P) \gamma V_{r^*,D} \\ &\leq \mu^P (1 - \gamma + \gamma V_{r^*}) + (1 - \mu^P) \gamma V_{r^*,D} \\ &= (1 - \gamma)\mu^P + \gamma V_{r^*} - (1 - \mu^P) \gamma (V_{r^*} - V_{r^*,D}) \\ &< (1 - \gamma)\mu^P + \gamma V_{r^*} - (1 - \gamma) (1 - \mu^P) g. \end{aligned}$$

Assuming that $\mu^P < 1$ (as must be the case for equilibria that are robust to $\varepsilon \neq 0$),

this is equivalent to

$$g < \frac{\mu^P - V_{r^*}}{1 - \mu^P}.$$

Finally, since $V_{r^*} \geq V_0 = (\mu^P)^2$, we have

$$g < \frac{\mu^P - (\mu^P)^2}{1 - \mu^P} = \mu^P.$$

5.3.2 Sufficient Conditions

As previewed in Section 1, we use different types of strategies to support cooperation when $g < 1$ and when $g \geq 1$ and $l > g + g^2$. The strategies used in the $g < 1$ case are similar to those used to prove Theorem 2(i): all records are either reciprocators (good-standing) or defectors (bad-standing), and players cycle between good and bad standing depending on the number of times they have been recorded as playing D (their “score”).

When $g \geq 1$, cooperation requires the presence of supercooperator records. Here we construct equilibria where players begin life as reciprocators; transition to being supercooperators once they have been recorded as playing D a certain number of times K ; and then permanently transition to being defectors once they have been recorded as playing D an additional L times, for some $K, L \in \mathbb{N}$. Two features of these equilibria are particularly notable.

First, players’ incentives to cooperate are provided solely by the threat of entering defector status once their score reaches $K + L$, at which point their reciprocator opponents switch from playing C to D against them. (In contrast, when a player switches from reciprocator to supercooperator status, their opponents’ behavior is unaffected.) Since the survival probability γ is less than 1, this threat looms larger the closer a player’s score is to $K + L$. Hence, players with higher scores are willing to incur greater costs to prevent their scores from increasing further. Our construction exploits this observation by finding a critical score K such that players with scores greater than K are willing to play C at a cost of l , while players with scores less than K are willing

to play C at a cost of g but not at a cost of l . That is, players with scores between K and L supercooperate, while those with scores less than K reciprocate.

Second, the feature that players with scores between K and L supercooperate rather than reciprocate may at first seem to make it harder to support cooperation: after all, defectors obtain higher payoffs against supercooperators than cooperators do. However, the presence of supercooperator records increases the steady-state share of reciprocators, via the following mechanism: Since players with scores between K and L supercooperate, their scores increase more slowly than if they reciprocated. Therefore, fewer players survive to enter defector status, which reduces the steady-state share of defectors. Finally, when there are fewer defectors, reciprocators defect less often, and hence their scores increase more slowly, which increases the steady-state share of reciprocators. In sum, the presence of supercooperator records reduces the steady-state share of defectors and increases the steady-state share of reciprocators, which enables steady-state cooperation.²⁸

5.4 Personal Public Randomization

Theorem 5 has two limitations: It falls short of completely characterizing when limit efficiency is attainable due to the presence of the “integer problem” term $f(g)$, and it only shows that efficiency is attainable for the iterated limit where first $\gamma \rightarrow 1$ and then $\varepsilon \rightarrow 0$, rather than the general double limit $(\gamma, \varepsilon) \rightarrow (1, 0)$.

Theorem 6 (proved in OA.5) shows that enriching the information structure to allow a simple class of randomizing devices, *personal public randomization* (PPR), circumvents the “integer problem” and also allows the construction of asymptotically efficient equilibria in the double limit. With PPR, each recorded action is associated

²⁸At the risk of over-interpreting our stylized model, we may relate this behavior to the “graduated sanctions” identified by Ostrom (1990) as essential to supporting cooperation in many communities. Ostrom found that graduated sanctions—where norm violators are given opportunities to resume cooperation before facing harsher punishments—help sustain cooperation by preventing excessively fast breakdowns following occasional violations. In our model, the mild punishment of transitioning to supercooperator status—at which point individuals increase their cooperation to prevent their status from deteriorating further—serves a broadly similar role.

in the player’s record with an indelible realization of an independent, uniform $[0, 1]$ random variable, which is observable to the player’s future partners. PPR is thus a “decentralized” form of public randomization, since randomizations attach to each player’s record independently.²⁹

The proof of Theorem 6(i) is the same as that of Theorem 5(i), and the $g < 1$ case of Theorem 6(ii) is similar to the corresponding case in Theorem 5(ii). The $l > g + g^2$ part of Theorem 6(ii) extends the “preciprocator→supercooperator→defector” strategies used to prove the corresponding case in Theorem 5(ii) in two ways. First, the transition rates between these phases are now tuned using the real-valued PPR realizations, rather than the discrete threshold scores K and L . This greater flexibility lets us eliminate the nuisance term f . Second, we add an additional defector phase at the beginning of each player’s lifetime. This gives us further flexibility to tune the population share of defectors, and in particular to keep it away from 0 when $\varepsilon_C \ll 1 - \gamma$. This is what lets us cover the double limit $(\gamma, \varepsilon) \rightarrow (1, 0)$ rather than only the iterated limit.

6 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 their opponents’ past play is summarized by noisy “records.” We study steady-state equilibria in a large population of players with geometrically distributed lifetimes, focusing on situations where there is little noise and lifetimes are long.

We find that any strictly individually rational outcome can be supported with second-order records, while with first-order records an outcome can be supported if it has a corresponding unprofitable punishment. In the prisoner’s dilemma, cooperation can be supported if and only if stage-game payoffs are either strictly supermodular and

²⁹We do not take a position on whether allowing or disallowing PPR is more realistic, but the two extreme cases of no correlating devices whatsoever and PPR give a sense of the range of outcomes that could arise for a variety of types of correlating devices.

mild or strongly supermodular. The strength of the short-term coordination motive and the temptation to cheat thus determine the prospects for robust long-term cooperation.

We conclude by discussing some possible extensions and alternative models.

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

Simpler strategies. It is also interesting to consider simpler classes of strategies. In Clark, Fudenberg, and Wolitzky (2020), we analyze the performance of tolerant grim trigger strategies in the PD, where players reciprocate until they have been recorded as playing D K times and defect thereafter. We find that when $g < l/(1 + l)$ such strategies can limit-support any cooperation share between g and $l/(1 + l)$, and that otherwise they cannot limit-support any positive cooperation share.

Sequential moves. In any strict equilibrium with first or second-order records, if players can “jump the gun” by taking their action before the opponent has a chance to respond, then only static equilibrium behavior can be supported.³⁰ However, our simultaneous-move specification applies not only when actions are literally simultaneous, but also whenever both players must choose their actions before fully observing

³⁰To see why, note that by jumping the gun a player can obtain stage-game payoff $\max_r u(s(r, r'), s(r', r))$ when matched with an opponent with record r' , by taking action $s(\arg \max_r u(s(r, r'), s(r', r)), r')$. This implies that all players must receive the same payoff when matched with each possible opponent.

their opponent's action. This seems like a natural reduced-form model for the typical case where cooperation unfolds gradually within each match.³¹

Multiple populations. It is easy to adapt our model to settings with multiple populations of players. Here efficient outcomes can always be fully supported in situations with one-sided incentive problems.³² For example, suppose a population of player 1's and a population of player 2's repeatedly play a product choice game, where only player 1 faces binding moral hazard at the efficient action profile (and player 2 wants to match player 1's action). The efficient outcome can always be supported with the following trigger strategies (with K chosen appropriately as a function of γ and ε): in each match, both partners play C if player 1's score is less than K , and both play D if player 1's score is greater than K .

Coarser record systems. In the PD, a simple noisy count of the number of times each player has defected enables cooperation to be iterated-limit-supported for all parameters that allow cooperation under any first-order record system (modulo the nuisance requirement $l > f(g)$). Coarser record systems cannot always do as well. For example, one natural information structure arises when records are "lost" with some probability each period, so a player's record is occasionally reset to 0. Our iterated-limit efficiency results easily extend to the triple iterated limit where first the resetting probability goes to 0, then $\gamma \rightarrow 1$, and then $\varepsilon \rightarrow 0$. The situation is different if a player can reset their record at will, for example by re-entering the game under a pseudonym. In this case, society must use strategies where new players must "build a reputation" before anyone will cooperate with them, as in Friedman and Resnick (2001).

Endogenous record systems. This paper has considered how features of an exogenously given record system (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

³¹If records are interdependent rather than second order, the strategies used to prove Theorem 2 remain equilibria for any possible move order.

³²Proposition 4 of Kandori (1992) is a similar result in a fixed-population model without noise.

observations, either to a central database or directly to other individual players. Intuitively, first-order information is relatively easy to extract, since if a player is asked to report only their partner’s behavior, they have no reason to lie as this information does not affect their own future record. Whether and how society can obtain higher-order information is an interesting question for future study.³³

References

- [1] Sachin Adlakha, Ramesh Johari, and Gabriel Y Weintraub. “Equilibria of Dynamic Games with Many Players: Existence, Approximation, and Market Structure”. *Journal of Economic Theory* (2015).
- [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 and Caroline Thomas. “Community Enforcement of Trust with Bounded Memory”. *Review of Economic Studies* (2018).
- [6] Samuel Bowles and Herbert Gintis. *A Cooperative Species: Human Reciprocity and its Evolution*. Princeton, NJ, US. Princeton University Press, 2011.
- [7] Robert Boyd et al. “The Evolution of Altruistic Punishment”. *Proceedings of the National Academy of Sciences* (2003).
- [8] Daniel Clark, Drew Fudenberg, and Alexander Wolitzky. *Indirect Reciprocity with Simple Records*. Tech. rep. 2020.
- [9] Daniel Clark, Drew Fudenberg, and Alexander Wolitzky. *Steady-State Equilibria in Anonymous Repeated Games, I: Trigger Strategies in General Stage Games*. Tech. rep. 2019.
- [10] Karen Clay. “Trade without Law: Private-order Institutions in Mexican California”. *The Journal of Law, Economics, and Organization* (1997).
- [11] Pedro Dal Bó and Guillaume R Fréchet. “On the Determinants of Cooperation in Infinitely Repeated Games: A Survey”. *Journal of Economic Literature* (2018).

³³Another possibility would be to consider “optimal” record systems subject to some constraints, in the spirit of information design.

- [12] Joyee Deb, Takuo Sugaya, and Alexander Wolitzky. “The Folk Theorem in Repeated Games with Anonymous Random Matching”. 2018.
- [13] Chrysanthos Dellarocas. “Reputation Mechanism Design in Online Trading Environments with Pure Moral Hazard”. *Information Systems Research* (2005).
- [14] Francesc Dilmé. “Helping Behavior in Large Societies”. *International Economic Review* (2016).
- [15] Robert C Ellickson. *Order without Law*. Harvard University Press, 1991.
- [16] Glenn Ellison. “Cooperation in the Prisoner’s Dilemma with Anonymous Random Matching”. *Review of Economic Studies* (1994).
- [17] Joseph Farrell and Eric Maskin. “Renegotiation in Repeated Games”. *Games and Economic Behavior* (1989).
- [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] Drew Fudenberg and Kevin He. “Learning and Type Compatibility in Signaling Games”. *Econometrica* (2018).
- [21] Drew Fudenberg and Kevin He. “Player-Compatible Equilibrium”. 2019.
- [22] Drew Fudenberg, David Levine, and Eric Maskin. “The Folk Theorem with Imperfect Public Information”. *Econometrica* (1994).
- [23] Drew Fudenberg, David G. Rand, and Anna Dreber. “Slow to Anger and Fast to Forgive: Cooperation in an Uncertain World”. *American Economic Review* (2012).
- [24] Herbert Gintis et al. “Explaining Altruistic Behavior in Humans”. *Evolution and Human Behavior* (2003).
- [25] Avner Greif. “Contract Enforceability and Economic Institutions in Early Trade: The Maghribi Traders’ Coalition”. *The American Economic Review* (1993).
- [26] Avner Greif, Paul Milgrom, and Barry R Weingast. “Coordination, Commitment, and Enforcement: The Case of the Merchant Guild”. *Journal of Political Economy* (1994).
- [27] Francesco Guala. “Reciprocity: Weak or Strong? What Punishment Experiments do (and do not) Demonstrate”. *Behavioral and Brain Sciences* (1 2012).
- [28] Yuval Heller and Erik Mohlin. “Observations on Cooperation”. *Review of Economic Studies* (2018).
- [29] Hugo A Hopenhayn. “Entry, Exit, and Firm Dynamics in Long Run Equilibrium”. *Econometrica* (1992).

- [30] Johannes Hörner and Wojciech Olszewski. “The Folk Theorem for Games with Private Almost-perfect Monitoring”. *Econometrica* (2006).
- [31] Boyan Jovanovic and Robert W Rosenthal. “Anonymous Sequential Games”. *Journal of Mathematical Economics* (1988).
- [32] Michihiro Kandori. “Social Norms and Community Enforcement”. *The Review of Economic Studies* (1992).
- [33] Nobuhiro Kiyotaki and Randall Wright. “A Search-theoretic Approach to Monetary Economics”. *The American Economic Review* (1993).
- [34] Benjamin Klein and Keith B Leffler. “The Role of Market Forces in Assuring Contractual Performance”. *Journal of Political Economy* (1981).
- [35] Daniel B Klein. “Promise Keeping in the Great Society: A Model of Credit Information Sharing”. *Economics & Politics* (1992).
- [36] Narayana Kocherlakota. “Money is Memory”. *Journal of Economic Theory* (1998).
- [37] 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).
- [38] Martin A. Nowak and Karl Sigmund. “Evolution of Indirect Reciprocity by Image Scoring”. *Nature* (1998).
- [39] Masahiro Okuno-Fujiwara and Andrew Postlewaite. “Social Norms and Random Matching Games”. *Games and Economic Behavior* (1995).
- [40] Elinor Ostrom. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge university press, 1990.
- [41] A Jorge Padilla and Marco Pagano. “Sharing Default Information as a Borrower Discipline Device”. *European Economic Review* (2000).
- [42] Karthik Panchanathan and Robert Boyd. “A Tale of Two Defectors: The Importance of Standing for Evolution of Indirect Reciprocity”. *Journal of Theoretical Biology* (2003).
- [43] 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.
- [44] Robert W. Rosenthal. “Sequences of Games with Varying Opponents”. *Econometrica* (1979).
- [45] Robert W. Rosenthal and Henry J. Landau. “A Game-theoretic Analysis of Bargaining with Reputations”. *Journal of Mathematical Psychology* (1979).
- [46] Robert Sugden. “New Developments in the Theory of Choice Under Uncertainty”. *Bulletin of Economic Research* (1986).

[47] 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 4. *There exists a family of finite subsets of R , $\{L(t, \eta)\}_{t \in \mathbb{N}, \eta > 0}$, such that*

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

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

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

player with record in $L(t-1, \eta/2)$ has a next-period record in $L(t, \eta)$ exceeds λ . For any $\mu \in \Delta(R)$, it then follows that $\sum_{r \in L(t-1, \eta/2)} \mu_r \geq (1 - \eta/2)(1 - \gamma)\gamma^{t-1}$ implies $\sum_{r \in L(t, \eta)} f_\sigma(\mu)[r] \geq \lambda(1 - \eta/2)(1 - \gamma)\gamma^{t-1} \geq (1 - \eta)(1 - \gamma)\gamma^t$. ■

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

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

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

Note that \bar{M} is convex and, by Corollary 4, must contain every steady-state distribution μ . The next lemma uses Corollary 4 to show that \bar{M} is non-empty.

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

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

For a given $t \in \mathbb{N}$, Corollary 4 implies that $\sum_{r \in L(t, \eta)} \mu_r^{i_j} \geq (1 - \eta)(1 - \gamma)\gamma^t$ for all $\eta > 0$ and all sufficiently high $j \in \mathbb{N}$, so

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

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

The following three claims imply that f_σ has a fixed point in \bar{M} ,³⁴ which completes the proof of parts 1 and 2 of Theorem 1.

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

Claim 2. f_σ maps \bar{M} to itself.

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

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

Here we show that $\lim_{j \rightarrow \infty} \mu^{i_j} = \tilde{\mu}$. For a given $\eta > 0$, there is a finite subset of records $L(\eta/2) \subset R$ such that $\sum_{r \in L(\eta/2)} \mu_r > 1 - \eta/2$ for every $\mu \in \bar{M}$. Thus, $|\mu_r^{i_j} - \tilde{\mu}_r| < \eta/2$ 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 4 imply that $\sum_{r \in L(t, \eta)} f_\sigma(\mu)[r] \geq (1 - \eta)(1 - \gamma)\gamma^t$ for all $t \in \mathbb{N}, \eta > 0$. Furthermore, $f_\sigma(\mu)[0] = 1 - \gamma$, and for all $t > 0$, $\gamma \sum_{r \in R(t-1)} \mu_r = \sum_{r \in R(t)} f_\sigma(\mu)[r]$, so $\sum_{r \in R(t-1)} \mu_r = (1 - \gamma)\gamma^{t-1}$ gives $\sum_{r \in R(t)} f_\sigma(\mu)[r] = (1 - \gamma)\gamma^t$. ■

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

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

³⁴This 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 Theorem 2

A.2.1 Proof of Theorem 2(i)

Let M be a positive integer such that $(u(a, a) - u(b, b))M > \max_{a'}\{\max\{u(a', a) - u(a, a), u(a', b) - u(b, b)\}\}$. We show that, with this choice of M , action a can be limit-supported by the cyclic strategies defined in Section 3, which we denote by σ^* .³⁵

Let $\tilde{\varepsilon}_{(a,a)} = \sum_{(\tilde{a}, \tilde{a}') \neq (a,a), (b,b)} \varepsilon_{(a,a), (\tilde{a}, \tilde{a}'})$ be the probability that the stage-game outcome is recorded as something other than (a, a) or (b, b) when the actual outcome is (a, a) , $\tilde{\varepsilon}_{(b,b)} = \sum_{(\tilde{a}, \tilde{a}') \neq (a,a), (b,b)} \varepsilon_{(b,b), (\tilde{a}, \tilde{a}'})$ be the probability that the outcome is recorded as something other than (a, a) or (b, b) when the actual outcome is (b, b) , and $\hat{\varepsilon}_{(b,b)} = \sum_{(\tilde{a}, \tilde{a}') \neq (b,b)} \varepsilon_{(b,b), (\tilde{a}, \tilde{a}'})$ be the probability that the outcome is recorded as something other than (b, b) when the actual outcome is (b, b) .

Consider a steady-state $\mu(\gamma, \varepsilon)$ for parameters (γ, ε) . Let $\mu^G(\gamma, \varepsilon)$ be the corresponding share of good-standing players. Similarly, for $i \in \{0, \dots, M-1\}$, let $\mu^{B_i}(\gamma, \varepsilon)$ be the share of bad-standing players who have accumulated i (b, b) profiles since last entering bad standing. We show that the unique limit point of any sequence of steady-state shares $(\mu^G(\gamma, \varepsilon), \mu^{B_0}(\gamma, \varepsilon), \dots, \mu^{B_{M-1}}(\gamma, \varepsilon))$ as $(\gamma, \varepsilon) \rightarrow (1, 0)$ is $(\tilde{\mu}^G, \tilde{\mu}^{B_0}, \dots, \tilde{\mu}^{B_{M-1}}) = (1, 0, \dots, 0)$. This implies that $\lim_{(\gamma, \varepsilon) \rightarrow (1, 0)} \mu^G(\gamma, \varepsilon) = 1$, so the share of good-standing players converges to 1 in the $(\gamma, \varepsilon) \rightarrow (1, 0)$ limit. Consequently, the population share of action a also converges to 1.

³⁵Note that the strategy σ^* does not depend on (γ, ε) .

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

We now show that $(\sigma^*, \mu(\gamma, \varepsilon))$ is a strict equilibrium when γ is sufficiently close to 1 and ε is sufficiently close to 0. For $0 \leq i < M - 1$, the value functions in the bad-standing phase B_i and the subsequent bad-standing phase B_{i+1} satisfy

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

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

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

Combining $\lim_{(\gamma, \varepsilon) \rightarrow (1, 0)} \mu^G(\gamma, \varepsilon) = 1$ with $V^G = \mu^G(\gamma, \varepsilon)^2 u(a, a) + (1 - \mu^G(\gamma, \varepsilon)^2)u(b, b)$ shows that $\lim_{(\gamma, \varepsilon) \rightarrow (1, 0)} V^G = u(a, a)$. Thus, taking the limits of these equations as $(\gamma, \varepsilon) \rightarrow (1, 0)$ gives $\lim_{(\gamma, \varepsilon) \rightarrow (1, 0)} V^{B_i} = \lim_{(\gamma, \varepsilon) \rightarrow (1, 0)} V^G = u(a, a)$ for all $i \in \{0, \dots, M - 1\}$.

A player in bad-standing phase i where $0 \leq i < M - 1$ strictly prefers to play b

against b when $(1 - \gamma)u(b, b) + \gamma\hat{\varepsilon}_{(b,b)}V^{B_i} + \gamma(1 - \hat{\varepsilon}_{(b,b)})V^{B_{i+1}} > (1 - \gamma)u(a', b) + \gamma(1 - \varepsilon_{(a',b),(b,b)})V^{B_i} + \gamma\varepsilon_{(a',b),(b,b)}V^{B_{i+1}}$ holds for $a' \neq b$. Manipulating this gives $(1 - \hat{\varepsilon}_{(b,b)} - \varepsilon_{(a',b),(b,b)})\gamma(V^{B_{i+1}} - V^{B_i})/(1 - \gamma) > u(a', b) - u(b, b)$. Equation 2 can be rewritten as

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

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

An almost identical argument shows that the incentives of players in bad-standing phase $M - 1$ are satisfied for (γ, ε) sufficiently close to $(1, 0)$. Thus, all that remains is to show that the incentives of players in good-standing are satisfied in the limit. A good-standing player has strict incentives to play a against a when $(1 - \gamma)u(a, a) + \gamma(1 - \tilde{\varepsilon}_{(a,a)})V^G + \gamma\tilde{\varepsilon}_{(a,a)}V^{B_0} > (1 - \gamma)u(a', a) + \gamma(\varepsilon_{(a',a),(a,a)} + \varepsilon_{(a',a),(b,b)})V^G + \gamma(1 - \varepsilon_{(a',a),(a,a)} - \varepsilon_{(a',a),(b,b)})V^{B_0}$ holds for $a' \neq a$. Manipulating this gives $(1 - \tilde{\varepsilon}_{(a,a)} - \varepsilon_{(a',a),(a,a)} - \varepsilon_{(a',a),(b,b)})\gamma(V^G - V^{B_0})/(1 - \gamma) > u(a', a) - u(a, a)$. Similarly, a good-standing player has strict incentives to play b against b when $(1 - \gamma)u(b, b) + \gamma(1 - \tilde{\varepsilon}_{(b,b)})V^G + \gamma\tilde{\varepsilon}_{(b,b)}V^{B_0} > (1 - \gamma)u(a', b) + \gamma(\varepsilon_{(a',b),(a,a)} + \varepsilon_{(a',b),(b,b)})V^G + \gamma(1 - \varepsilon_{(a',b),(a,a)} - \varepsilon_{(a',b),(b,b)})V^{B_0}$ holds for $a' \neq b$. Manipulating this gives $(1 - \tilde{\varepsilon}_{(b,b)} - \varepsilon_{(a',a),(a,a)} - \varepsilon_{(a',a),(b,b)})\gamma(V^G - V^{B_0})/(1 - \gamma) > u(a', b) - u(b, b)$. Combining Equations 2 and 3 gives

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

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

A.2.2 Proof of Theorem 2(ii)

We show that action a can be iterated-limit-supported by the tolerant grim trigger strategies defined in Section 3, with K chosen appropriately as a function of γ and ε . We then show that the resulting equilibria are coordination-proof in the iterated limit.

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

For every $k \in \{0, \dots, K-1\}$, let μ_k be the share of players with score k and let $\mu^G = \sum_{k=0}^{K-1} \mu_k$ be the total share of players in good standing. The inflow into score 0 is simply $1 - \gamma$, the share of newborn players. The outflow from score 0 is the sum of $(1 - \gamma)\mu_0$, the share of players with score 0 who die in a given period, and $\gamma(\tilde{\varepsilon}_{a,a}\mu^G + \tilde{\varepsilon}_{(b,b)}(1 - \mu^G))\mu_0$, the share of players who move from score 0 to score 1 in a given period. In a steady state, the inflow and outflow of score 0 must be equal, so we obtain $\mu_0 = (1 - \gamma)/(1 - \gamma + \gamma(\tilde{\varepsilon}_{(a,a)}\mu^G + \tilde{\varepsilon}_{(b,b)}(1 - \mu^G)))$. Similarly, for any $0 < k < K$, the inflow is $\gamma(\tilde{\varepsilon}_{(a,a)}\mu^G + \tilde{\varepsilon}_{(b,b)}(1 - \mu^G))\mu_{k-1}$, the share of players who move from score $k-1$ to score k . The outflow from score k is the sum of $(1 - \gamma)\mu_k$, the share of players with score k who die, and $\gamma(\tilde{\varepsilon}_{(a,a)}\mu^G + \tilde{\varepsilon}_{(b,b)}(1 - \mu^G))\mu_k$, the share of players who move from score k to score $k+1$. Thus, in steady state, $\mu_k = \gamma(\tilde{\varepsilon}_{(a,a)}\mu^G + \tilde{\varepsilon}_{(b,b)}(1 - \mu^G))\mu_{k-1}/(1 - \gamma + \gamma(\tilde{\varepsilon}_{(a,a)}\mu^G + \tilde{\varepsilon}_{(b,b)}(1 - \mu^G)))$. Combining these results gives $\mu_k = \xi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)^k(1 - \xi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G))$ for every $k \in \{0, \dots, K-1\}$, where $\xi : [0, 1] \times (0, 1) \times (0, 1) \times [0, 1] \rightarrow [0, 1]$ is the function given by

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

(Intuitively, $\xi(\cdot)$ gives the probability that a good-standing player lives to see their score increase by 1, as a function of the survival probability, noise levels, and good-standing population share.)

Summing $\mu_k = \xi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)^k(1 - \xi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G))$ over $k \in \{0, \dots, K-1\}$

gives

$$\mu^G = 1 - \xi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)^K. \quad (4)$$

It follows that there is a steady state with parameters (γ, ε) , tolerance level K , and share of good-standing players μ^G precisely when Equation 4 holds. We use this in OA.1 to show that, to arbitrary precision, any good-standing population share $\mu^G \in [0, 1]$ can be achieved in a steady state for fixed ε in the limit as $\gamma \rightarrow 1$ by appropriately tailoring K . Formally, we prove the following lemma.

Lemma 6. *Fix ε . For all $\Delta > 0$, there exists $\bar{\gamma} < 1$ such that, for all $\gamma > \bar{\gamma}$ and $\mu^G \in [0, 1]$, there exists a $\hat{\mu}^G$ with $|\hat{\mu}^G - \mu^G| < \Delta$ that satisfies Equation 4 for some K .*

Now, we turn to the incentives of the good-standing players.³⁶ The expected payoff of a good-standing player with score k from playing a against a is $(1 - \gamma)u(a, a) + \gamma(1 - \tilde{\varepsilon}_{a,a})V_k + \gamma\tilde{\varepsilon}_{(a,a)}V_{k+1}$, while their expected payoff from playing some action $a' \neq a$ is $(1 - \gamma)u(a', a) + \gamma(\varepsilon_{(a',a),(a,a)} + \varepsilon_{(a',a),(b,b)})V_k + \gamma(1 - \varepsilon_{(a',a),(a,a)} - \varepsilon_{(a',a),(b,b)})V_{k+1}$. Such a player will prefer to play a against a rather than playing any other action a' if $(1 - \tilde{\varepsilon}_{(a,a)} - \varepsilon_{(a',a),(a,a)} - \varepsilon_{(a',a),(b,b)})\gamma(V_k - V_{k+1})/(1 - \gamma) > u(a', a) - u(a, a)$ for all $a' \neq a$. Similarly, the incentives of good-standing player with score k to play b against b are satisfied if $(1 - \tilde{\varepsilon}_{(b,b)} - \varepsilon_{(a',b),(a,a)} - \varepsilon_{(a',b),(b,b)})\gamma(V_k - V_{k+1})/(1 - \gamma) > u(a', b) - u(b, b)$ for all $a' \neq b$. Thus, it follows that the incentives of good-standing players both to play a against a and play b against b are satisfied if $\gamma(V_k - V_{k+1})/(1 - \gamma)$ is sufficiently high.

For $k < K$, the relationship between the value function at score k and the value function at score $k + 1$ is $V_k = (1 - \gamma)(\mu^G u(a, a) + (1 - \mu^G)u(b, b)) + \gamma(1 - \tilde{\varepsilon}_{(a,a)}\mu^G - \tilde{\varepsilon}_{(b,b)}(1 - \mu^G))V_k + \gamma(\tilde{\varepsilon}_{(a,a)}\mu^G + \tilde{\varepsilon}_{(b,b)}(1 - \mu^G))V_{k+1}$. Combining this with the fact that the bad-standing value functions are $V_k = u(b, b)$ for $k \geq K$, it follows that

$$\begin{aligned} V_k = & (1 - \xi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)^{K-k})(\mu^G u(a, a) + (1 - \mu^G)u(b, b)) \\ & + \xi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)^{K-k}u(b, b) \end{aligned} \quad (5)$$

³⁶Since (b, b) is a strict static equilibrium, bad-standing players' incentives are always satisfied.

for $k < K$. An implication of Equation 5 is that $V_0 - V_1 \leq V_k - V_{k+1}$ for all $k < K$, so the incentives of good-standing players with score 0 will be the most difficult to satisfy. Moreover, combining Equations 4 and 5 gives

$$\frac{\gamma}{1 - \gamma}(V_0 - V_1) = \frac{\mu^G(1 - \mu^G)}{\tilde{\varepsilon}_{(a,a)}\mu^G + \tilde{\varepsilon}_{(b,b)}(1 - \mu^G)}(u(a, a) - u(b, b)). \quad (6)$$

Note that $\tilde{\varepsilon}_{(a,a)}\mu^G + \tilde{\varepsilon}_{(b,b)}(1 - \mu^G)$ becomes arbitrarily small as ε approaches 0. Thus, Equation 6 implies that, for a sequence of steady states in which the good-standing population share remains bounded away from 0 and 1, the incentives of good-standing players will be satisfied as $\varepsilon \rightarrow 0$. Combining this fact with Lemma 6, we conclude that (a, a) can be iterated-limit-supported by strict equilibria.

A.2.2.2 Proof of Coordination-Proofness

We first argue that in every match between bad-standing players, there is no Nash equilibrium in the augmented game that Pareto-dominates (b, b) . This is because once a player enters bad standing, their value function remains constant at $u(b, b)$. Thus, any Nash equilibrium in the augmented game between two bad-standing players must also be a static equilibrium in the stage game. Since there is no static equilibrium that Pareto-dominates (b, b) , it follows that two bad-standing players playing (b, b) is coordination-proof.

Now we show that in any match involving a good-standing player, there is no Nash equilibrium in the augmented game that Pareto-dominates the action profile the players are supposed to play. Equation 6 shows that $\gamma(V_0 - V_1)/(1 - \gamma) \rightarrow \infty$ as $\varepsilon \rightarrow 0$ for a sequence of steady states in which the good-standing population share remains bounded away from 0 and 1. Not only does this imply that the incentive constraints of good-standing players are satisfied along such a sequence when ε is sufficiently close to 0, it also implies that no good-standing player would ever prefer an action profile other than (a, a) or (b, b) be played in one of their matches. Thus, in any match involving a good-standing player, we need only consider whether (a, a) or (b, b)

are Nash equilibria in the augmented game and whether one of these profiles Pareto-dominates the other. When two good-standing players match, both (a, a) and (b, b) are Nash equilibria in the augmented game, but (b, b) does not Pareto-dominate (a, a) . Indeed, if (b, b) did Pareto-dominate (a, a) , this would imply that the value functions for these good-standing players would be no higher than $u(b, b)$, which is not possible given that $u(a, a) > u(b, b)$. Thus, the prescribed play between two good-standing players is coordination-proof. Moreover, in any match involving a bad-standing player, all Nash equilibria in the augmented game require the bad-standing player to play a static best-response to the action of their opponent. Given that (a, a) is not a static equilibrium (since $u(a, a) > u(b, b)$ and (b, b) is not Pareto-dominated by any static equilibrium), we conclude that (b, b) is coordination-proof in a match between a good-standing player and a bad-standing player.

A.3 Proof of Theorem 5

The proof proceeds as follows. Section A.3.1 derives the incentive constraints that must be satisfied in any strict, coordination-proof equilibrium with noisy first-order records. Section A.3.2 proves Theorem 5(i) (necessary conditions for cooperation). The main step is proving Lemma 10, which shows that $\mu^P + \mu^S(l - g) > g$ in any strict, coordination-proof equilibrium with $\mu^C > 0$. Section A.3.3 proves Theorem 5(ii) (sufficient conditions for cooperation). This is the longest part of the proof and is split into three parts: Section A.3.3.1 shows that cyclic strategies can iterated-limit-support cooperation when $g < 1$; Section A.3.3.2 shows that “preciprocator \rightarrow supercooperator \rightarrow defector” strategies can iterated-limit-support cooperation when $l > \max\{g + g^2, f(g)\}$; and Section A.3.3.3 shows that both of these classes of strategies are coordination-proof.

A.3.1 Incentive Constraints with Noisy Records

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

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

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

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

Lemma 7. *For any noisy record system,*

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

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

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

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

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

The strategies we use to prove part (ii) of the theorem depend on a player's record only through the number of times they have been recorded as playing D , their *score*. For such *scoring strategies*, we slightly abuse notation in writing V_k for the continuation payoff of a player with score k .³⁷ The incentive constraints take a simpler form with such strategies: For all k we have $\varepsilon_C(k) = \varepsilon_C$, $\varepsilon_D(k) = \varepsilon_D$, $V_k^C = V_k$, and $V_k^D = V_{k+1}$. The normalized reward thus simplifies to

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

Lemma 8. *For scoring strategies, Lemma 7 holds with $W_k = (1 - \varepsilon_C - \varepsilon_D)(\pi_k - V_k)/p_k^D$.*

A.3.2 Proof of Theorem 5(i)

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

Lemma 9. *For any first-order record system, in any strict equilibrium, $\mu^S < 1/(1 + g)$.*

Lemma 10. *For any noisy first-order record system, in any strict, coordination-proof equilibrium with $\mu^C > 0$, $\mu^P + \mu^S(l - g) > g$.*

³⁷Recall that V_r is defined as the continuation value of a player with record r . Under scoring strategies, two players with different records that share the same score have the same continuation value, so we can index V by k rather than r .

Lemma 9 says that there cannot be too many supercooperators. It holds because new players with record 0 have the option of always playing D , so in any strict equilibrium with $\mu^C > 0$, it must be that $\mu^S(1+g) < V_0 \leq 1$, which gives $\mu^S < 1/(1+g)$.

Conversely, Lemma 10 implies that cooperation requires a positive share of supercooperators when $g \geq 1$, and moreover that the required share grows when g and l are increased by the same amount. It is proved in the next subsection.

Theorem 5(i) follows from Lemmas 9 and 10 since, by Lemma 3, it is impossible to satisfy both $\mu^S < 1/(1+g)$ and $\mu^P + \mu^S(l-g) > g$ when $g \geq 1$ and $l \leq g+g^2$.

A.3.2.1 Necessary Conditions for Cooperation and Proof of Lemma 10

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

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

Proof. Suppose otherwise. Since $V_r = (1-\gamma)\pi_r + \gamma(1-p_r^D)V_r^C + \gamma p_r^D V_r^D$ and $\pi_{r_n} = \mu^S(1+g)$ for all r_n , we have $V_{r_n} = (1-\gamma)\mu^S(1+g) + \gamma(1-p_{r_n}^D)V_{r_n}^C + \gamma p_{r_n}^D V_{r_n}^D$ for all r_n . This implies

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

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

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

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

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

Now, consider the case where $\bar{V} > \mu^P \mu^C + \mu^S(\mu^C - \mu^D l) + \mu^D \mu^S(1+g)$. For any sequence of records $\{r_n\}_{n \in \mathbb{N}}$ such that $\lim_{n \rightarrow \infty} V_{r_n} = \bar{V}$, $\lim_{n \rightarrow \infty} \max\{V_{r_n}^C - V_{r_n}, 0\} = 0$, so there is some sufficiently high n such that $V_{r_n} - (\gamma/(1-\gamma))(V_{r_n}^C - V_{r_n}) \geq \mu^P \mu^C + \mu^S(\mu^C - \mu^D l) + \mu^D \mu^S(1+g)$. Additionally, by Lemma 11, for sufficiently high n , the record r_n must be either a preciprocator or a supercooperator. \blacksquare

We are now ready to prove Lemma 10.

Proof of Lemma 10. First, take the case where r' is a preciprocator. Then by Lemma 7, we must have

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

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

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

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

Now, take the case where r' is a supercooperator. Then by Lemma 7, we must have $\pi_{r'} - V_{r'} + (\gamma/(1-\gamma))(V_{r'}^C - V_{r'}) > 0$. When $\pi_{r'} = \mu^C - \mu^D l$ and $V_{r'} - \frac{\gamma}{1-\gamma}(V_{r'}^C - V_{r'}) \geq \mu^P \mu^C + \mu^S(\mu^C - \mu^D l) + \mu^D \mu^S(1+g)$, this implies

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

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

A.3.3 Proof of Theorem 5(ii)

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

We show that a class of “ $(P_K D_M)^\infty$ ” cyclic strategies where K and M are positive integers can limit-support C when $g < 1$. The strategy $(P_K D_M)^\infty$ is such that a player is a reciprocator if their score $k \bmod K+M$ is strictly less than K and a defector if $k \bmod K+M$ weakly exceeds K . We refer to a player’s score $\bmod K+M$ as their *phase* and denote it with i . Thus, in the $(P_K D_M)^\infty$ strategy, a player is a reciprocator when $i < K$ and a defector when $K \leq i < K+M$.

Fix any rational number ρ satisfying $g < \rho < \min\{l, 1\}$. Let K and M be integers such that $K \geq M > 0$ and $M/K = \rho$. We show here that $(P_K D_M)^\infty$ limit-supports cooperation; the proof that these equilibria are coordination proof is in A.3.3.3.

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

$$\bar{\mu}^C(\varepsilon) = \frac{(1 - \varepsilon_D)(1 + \rho) - \sqrt{(1 - \varepsilon_D)^2(1 + \rho)^2 - 4(1 - \varepsilon_D)(1 - \varepsilon_C - \varepsilon_D)\rho}}{2(1 - \varepsilon_C - \varepsilon_D)\rho}. \quad (7)$$

Lemma 13. *With cyclic strategies, let $\mu^C(\gamma, \varepsilon)$ denote the share of cooperators in some steady state for arbitrary $(\gamma, \varepsilon) \in (0, 1) \times (0, 1) \times (0, 1)$. Then $\lim_{\gamma \rightarrow 1} \mu^C(\gamma, \varepsilon) = \bar{\mu}^C(\varepsilon)$.*

Proof of Lemma 13. Let $\mu_i(\gamma, \varepsilon)$ be the share of players in phase i for some steady state at parameters (γ, ε) , and similarly let $\mu^C(\gamma, \varepsilon) = \sum_{i=0}^{K-1} \mu_i(\gamma, \varepsilon)$ be the corre-

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

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

$$(1 - \varepsilon_D - (1 - \varepsilon_C - \varepsilon_D)\tilde{\mu}^C(\varepsilon))\tilde{\mu}_{K-1}(\varepsilon) = (1 - \varepsilon_D)\tilde{\mu}_{K+M-1}(\varepsilon). \quad (8)$$

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

$$\begin{aligned} \mu_i(\gamma_j, \varepsilon) &= \frac{\gamma_j(1 - \varepsilon_D - (1 - \varepsilon_C - \varepsilon_D)\mu^C(\gamma_j, \varepsilon))}{1 - \gamma(\varepsilon_D + (1 - \varepsilon_C - \varepsilon_D)\mu^C(\gamma_j, \varepsilon))} \mu_{i-1}(\gamma, \varepsilon) \\ &= \beta(\gamma_j, \varepsilon, \mu^C)\mu_{i-1}(\gamma_j, \varepsilon), \end{aligned}$$

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

$$\tilde{\mu}_{K-1}(\varepsilon) = \frac{1}{K}\tilde{\mu}^C(\varepsilon). \quad (9)$$

When $K < i \leq K + M - 1$, both phase i and phase $i - 1$ are defectors. Therefore, the outflow from phase i is $(1 - \gamma\varepsilon_D)\mu_i(\gamma, \varepsilon)$, while the inflow into phase i is $\gamma(1 -$

$\varepsilon_D)\mu_{i-1}(\gamma, \varepsilon)$. Thus, $\mu_i(\gamma_j, \varepsilon) = \gamma_j(1 - \varepsilon_D)/(1 - \gamma\varepsilon_D)\mu_{i-1}(\gamma, \varepsilon)$, which consequently gives $\tilde{\mu}_i(\varepsilon) = \tilde{\mu}_{i-1}(\varepsilon)$. Since this holds for all $K < i \leq K + M - 1$, we conclude that $\tilde{\mu}_i(\varepsilon) = \tilde{\mu}_{K+M-1}(\varepsilon)$ for all $K \leq i \leq K + M - 1$, so

$$\tilde{\mu}_{K+M-1}(\varepsilon) = \frac{1}{M}(1 - \tilde{\mu}^C(\varepsilon)). \quad (10)$$

By Equations 8, 9, and 10, we obtain

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

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

The next step is to show that these strategies yield strict equilibria in the iterated limit. To do so, we prove the following lemma, which when combined with Lemma 8 and the fact that $p_k^D = 1 - \varepsilon_D - (1 - \varepsilon_C - \varepsilon_D)\mu^C$ for any preciprocator record, shows that the $(P_K D_M)^\infty$ strategy yield a strict equilibrium in the iterated limit.

Lemma 14. *Let $\mu^C(\gamma, \varepsilon)$ denote the share of cooperators and let $V_i(\gamma, \varepsilon)$ denote the value function for phase i in some steady state for arbitrary $(\gamma, \varepsilon) \in (0, 1) \times (0, 1) \times (0, 1)$.*

Then,

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

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

Proof of Lemma 14. First, note that for all $0 \leq i < K - 1$,

$$\begin{aligned} V_i(\gamma, \varepsilon) &= (1 - \gamma)\mu^C(\gamma, \varepsilon) + \gamma(\varepsilon_D + (1 - \varepsilon_C - \varepsilon_D)\mu^C(\gamma, \varepsilon))V_i(\gamma, \varepsilon) \\ &\quad + \gamma(1 - \varepsilon_D - (1 - \varepsilon_C - \varepsilon_D)\mu^C(\gamma, \varepsilon))V_{i+1}(\gamma, \varepsilon). \end{aligned}$$

Taking the limit as $\gamma \rightarrow 1$ gives $\lim_{\gamma \rightarrow 1} V_i(\gamma, \varepsilon) = \lim_{\gamma \rightarrow 1} V_{i+1}(\gamma, \varepsilon)$, should the limits exist. Since this holds for all $0 \leq i < K - 1$ and, by Lemma 13, $V_0(\gamma, \varepsilon) = (\mu^C(\gamma, \varepsilon))^2$ implies that $\lim_{\gamma \rightarrow 1} V_0(\gamma, \varepsilon) = (\bar{\mu}^C(\varepsilon))^2$, we conclude that $\lim_{\gamma \rightarrow 1} V_i(\gamma, \varepsilon) = (\bar{\mu}^C(\varepsilon))^2$ for all $0 \leq i \leq K - 1$.

Thus, all that remains to be shown is that

$$\lim_{\varepsilon \rightarrow 0} \frac{1 - \varepsilon_C - \varepsilon_D}{1 - \varepsilon_D - (1 - \varepsilon_C - \varepsilon_D)\bar{\mu}^C(\varepsilon)} (\bar{\mu}^C(\varepsilon) - (\bar{\mu}^C(\varepsilon))^2) = \rho. \quad (11)$$

Note that

$$\begin{aligned} & \frac{1 - \varepsilon_C - \varepsilon_D}{1 - \varepsilon_D - (1 - \varepsilon_C - \varepsilon_D)\bar{\mu}^C(\varepsilon)} (\bar{\mu}^C(\varepsilon) - (\bar{\mu}^C(\varepsilon_C \varepsilon_D))^2) \\ &= (1 - \varepsilon_C - \varepsilon_D) \left(\frac{\frac{1 - \bar{\mu}^C(\varepsilon)}{\varepsilon_C}}{\bar{\mu}^C(\varepsilon) + (1 - \varepsilon_D)\frac{1 - \bar{\mu}^C(\varepsilon)}{\varepsilon_C}} \right) \bar{\mu}^C(\varepsilon). \end{aligned}$$

Since $\lim_{\varepsilon \rightarrow 0} \bar{\mu}^C(\varepsilon) = 1$ and $\lim_{\varepsilon \rightarrow 0} (1 - \bar{\mu}^C(\varepsilon))/\varepsilon_C = \rho/(1 - \rho)$, as can be readily confirmed, we conclude that Equation 11 holds. \blacksquare

Combining Lemmas 13 and 14 shows C can be limit-supported when $g < 1$.

A.3.3.2 Limit-Supporting C when $l > \max\{g + g^2, f(g)\}$

We show that a class of “ $P_K S_L D_\infty$ ” strategies where K and L are positive integers can limit-support C when $l > \max\{g + g^2, f(g)\}$. The strategy $P_K S_L D_\infty$ is such that a player with score k is a preciprocator if $k < K$, a supercooperator if $K \leq k < K + L$, and a defector if $k \geq K + L$. We show here that K and L can be chosen as functions of γ , ε_C , and ε_D so that these strategies limit-support C . In A.3.3.3, we show that these equilibria are coordination proof.

Lemma 15. *When $l > \max\{g + g^2, f(g)\}$, C can be limit-supported by $P_K S_L D_\infty$ strategies.*

Note that when $g + g^2 > f(g)$ (i.e., $g > \approx 2.858$), limit efficiency is achieved if $l > g + g^2$. For smaller values of g , we use the stronger requirement that $l > f(g)$ to guarantee limit efficiency. We do not know whether the condition $l > f(g)$ is necessary; it comes from the combination of the record $K - 1$ preciprocator and record K supercooperator facing identical incentives as $\gamma \rightarrow 1$ and the fact that K and L must be integers.

To prove Lemma 15, we use the following lemma, which characterizes the population shares that are possible for given (γ, ε) with $P_K S_L D_\infty$ equilibria. Let $\alpha : (0, 1) \times (0, 1) \rightarrow (0, 1)$ and $\beta : (0, 1) \times (0, 1) \times [0, 1] \rightarrow (0, 1)$ be the functions defined by

$$\begin{aligned}\alpha(\gamma, \varepsilon_C) &= \frac{\gamma \varepsilon_C}{1 - \gamma(1 - \varepsilon_C)}, \\ \beta(\gamma, \varepsilon, \mu) &= \frac{\gamma(1 - \varepsilon_D - (1 - \varepsilon_C - \varepsilon_D)\mu)}{1 - \gamma(\varepsilon_D + (1 - \varepsilon_C - \varepsilon_D)\mu)}.\end{aligned}\tag{12}$$

Lemma 16. *There is a $P_K S_L D_\infty$ equilibrium with share of cooperators μ^C , share of reciprocators μ^P , and share of supercooperators μ^S if and only if the following conditions hold:*

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

$$\begin{aligned}2. \text{ Incentives:} \quad & (C|C)_0 : \frac{(1 - \varepsilon_C - \varepsilon_D)(1 - \mu^C)}{1 - \varepsilon_D - (1 - \varepsilon_C - \varepsilon_D)\mu^C} (\mu^P + \mu^S(l - g)) > g, \\ & (D|D)_{K-1} : \frac{\gamma(1 - \varepsilon_C - \varepsilon_D)(1 - \mu^C)}{1 - \gamma(\varepsilon_D + (1 - \varepsilon_C - \varepsilon_D)\mu^C)} (\mu^P + \mu^S(l - g)) + \mu^P l < l, \\ & (C|D)_K \text{ (if } \mu^S > 0) : \frac{(1 - \varepsilon_C - \varepsilon_D)(1 - \mu^C)}{1 - \varepsilon_D - (1 - \varepsilon_C - \varepsilon_D)\mu^C} (\mu^P + \mu^S(l - g)) + \mu^P l > l.\end{aligned}$$

The proof of Lemma 16 is in OA.3. The feasibility constraints come from calculating the steady-state shares μ_k for the strategy $P_K S_L D_\infty$ as a function of μ^C , and then setting $\mu^C = \sum_{k=0}^{K+L-1} \mu_k$, $\mu^P = \sum_{k=0}^{K-1} \mu_k$, and $\mu^S = \sum_{k=K}^{K+L-1} \mu_k$. The $(C|C)_0$ and $(C|D)_K$ incentive constraints come from solving V_0 and V_K , the value functions at the corresponding records, and using Lemma 8. The $(D|D)_{K-1}$ constraint is derived by using the value of V_K , after relating V_{K-1} and V_K .

We now identify a target level of cooperation $h(\varepsilon, \mu^P)$ for fixed ε and a certain range of feasible μ^P . We then show (i) that we can send the level of cooperation to

1 as $\varepsilon \rightarrow 0$, and (ii) that there are feasible profiles satisfying the incentive constraints where the level of cooperation actually attains this target as $\gamma \rightarrow 1$.

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

$$\frac{(1 - \varepsilon_C - \varepsilon_D)(1 - \mu^C)}{1 - \varepsilon_D - (1 - \varepsilon_C - \varepsilon_D)\mu^C} ((l - g)\mu^C + (1 + g - l)\mu^P) + l\mu^P = l \quad (13)$$

and the function $h(\varepsilon, \mu^P)$ defined by $h(\varepsilon, \mu^P) = \max\{\mu^C \in [0, 1] : \mu^C \text{ solves Equation 13}\}$. If $h(\varepsilon, \mu^P)$ is well-defined, it gives the maximum level of cooperation for the given μ^P , ε_C , and ε_D that satisfies the $\gamma \rightarrow 1$ “limit” incentive constraints of Lemma 16.

Lemma 17. *For any $\mu^P \in (g/(1+g), 1 - g/l]$, $h(\varepsilon, \mu^P)$ is well-defined for sufficiently small and positive $\varepsilon_C, \varepsilon_D$. Moreover, $\lim_{\varepsilon \rightarrow 0} h(\varepsilon, \mu^P) = 1$.*

Proof of Lemma 17. Straightforward calculations show that $h(\varepsilon, \mu^P)$ is well-defined for sufficiently small and positive $\varepsilon_C, \varepsilon_D$, and that

$$\lim_{\varepsilon \rightarrow 0} \frac{1 - h(\varepsilon, \mu^P)}{\varepsilon_C} = \frac{l(1 - \mu^P)}{(1 + g)\mu^P - g}. \quad (14)$$

An immediate implication of this is $\lim_{\varepsilon \rightarrow 0} h(\varepsilon, \mu^P) = 1$. ■

To complete the proof of Lemma 15, we combine $\lim_{\varepsilon \rightarrow 0} h(\varepsilon, \mu^P) = 1$ with the following lemma. Let $\bar{\mu}_{PSD}^C(\gamma, \varepsilon)$ be the maximal share of cooperators in any equilibrium using any $P_K S_L D_\infty$ generalized trigger strategy for parameters γ, ε_C , and ε_D :

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

Lemma 18. *If $l > \max\{g + g^2, f(g)\}$, there is some $\bar{\varepsilon} > 0$ and $\mu^P \in (g/(1+g), 1 - g/l]$ such that $\liminf_{\gamma \rightarrow 1} \bar{\mu}_{PSD}^C(\gamma, \varepsilon) \geq h(\varepsilon, \mu^P)$ for $\tilde{\varepsilon}_C, \tilde{\varepsilon}_D < \bar{\varepsilon}$.*

OA.4 presents the proof of Lemma 18. It first identifies a condition on μ^P that, if satisfied, enables the use of the inverse function theorem and other tools of differential calculus to show that, for sufficiently small ε_C and ε_D , any neighborhood of

$(h(\varepsilon, \mu^P), \mu^P)$ can be approached by feasible profiles for sufficiently high γ . It then shows that this condition is satisfied for some $\mu^P \in (g/(1+g), 1-g/l]$ whenever $l > \max\{g+g^2, f(g)\}$.

A.3.3.3 Proof of Coordination-Proofness

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

Online Appendix for “Steady-State Equilibria in Anonymous Repeated Games”

Daniel Clark, Drew Fudenberg, and Alexander Wolitzky

March 10, 2020

OA.1 Proof of Lemma 6

Lemma 6. *Fix ε . For all $\Delta > 0$, there exists $\bar{\gamma} < 1$ such that, for all $\gamma > \bar{\gamma}$ and $\mu^G \in [0, 1]$, there exists a $\hat{\mu}^G$ with $|\hat{\mu}^G - \mu^G| < \Delta$ that satisfies Equation 4 for some K .*

To prove Lemma 6, we calculate the number $\tilde{K} \in \mathbb{R}_+$ that satisfies $\mu^G = 1 - \xi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)^{\tilde{K}}$ as a function of $(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)$. Whenever \tilde{K} is an integer, it corresponds to a tolerance level at which the corresponding tolerant grim trigger strategy sustains μ^G in a steady state for parameters $(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)})$. We establish that, for fixed $(\tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)})$, the magnitude of the μ^G -partial derivative of \tilde{K} becomes arbitrarily large for almost all $\mu^G \in (0, 1)$ as $\gamma \rightarrow 1$. This implies that, for fixed $(\tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)})$, the distance between any $\mu^G \in [0, 1]$ and the set of feasible steady-state good-standing shares becomes arbitrarily small as $\gamma \rightarrow 1$.

*This paper replaces our earlier papers “Steady-State Equilibria in Anonymous Repeated Games, I: Trigger Strategies in General Stage Games,” “Steady-State Equilibria in Anonymous Repeated Games, II: Coordination-Proof Strategies in the Prisoner’s Dilemma,” and “Robust Cooperation with First-Order Information.” We thank Nageeb Ali, V Bhaskar, Glenn Ellison, Sander Heinsalu, Yuval Heller, Takuo Sugaya, Satoru Takahashi, and Caroline Thomas for helpful comments and conversations, and NSF grants SES 1643517 and 1555071 and Sloan Foundation grant 2017-9633 for financial support.

Let $\tilde{K} : (0, 1) \times (0, 1) \times (0, 1) \times (0, 1) \rightarrow \mathbb{R}_+$ be the function given by

$$\tilde{K}(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G) = \frac{\ln(1 - \mu^G)}{\ln(\xi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G))}. \quad (\text{OA } 1)$$

By construction, $\tilde{K}(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)$ is the unique $K \in \mathbb{R}_+$ such that $\mu^G = 1 - \xi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)^K$. Let $\psi : (0, 1] \times (0, 1) \times (0, 1) \times (0, 1) \rightarrow \mathbb{R}$ be the function given by

$$\psi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G) = \begin{cases} 1 + \ln(1 - \mu^G)(1 - \mu^G) \frac{\frac{\partial \xi}{\partial \mu^G}(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)}{\xi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G) \ln(\xi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G))} & \text{if } \gamma < 1 \\ 1 + \frac{(\tilde{\varepsilon}_{(b,b)} - \tilde{\varepsilon}_{(a,a)}) \ln(1 - \mu^G)(1 - \mu^G)}{\tilde{\varepsilon}_{(a,a)} \mu^G + \tilde{\varepsilon}_{(b,b)}(1 - \mu^G)} & \text{if } \gamma = 1 \end{cases}.$$

The μ^G derivative of $\tilde{K}(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)$ is related to $\psi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)$ in the following lemma.

Lemma OA 1. $\tilde{K} : (0, 1) \times (0, 1) \times (0, 1) \times (0, 1) \rightarrow \mathbb{R}_+$ is differentiable in μ^G with derivative given by

$$\frac{\partial \tilde{K}}{\partial \mu^G}(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G) = - \frac{\psi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)}{(1 - \mu^G) \ln(\xi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G))}.$$

Proof of Lemma OA 1. From Equation OA 1, it follows that $\tilde{K}(\gamma, \varepsilon, \mu^C)$ is differentiable in μ^C with derivative given by

$$\begin{aligned} \frac{\partial \tilde{K}}{\partial \mu^G}(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G) &= - \frac{\frac{\ln(\xi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G))}{1 - \mu^G} + \frac{\ln(1 - \mu^G) \frac{\partial \xi}{\partial \mu^G}(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)}{\xi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)}}{\ln(\xi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G))^2} \\ &= - \frac{1 + \ln(1 - \mu^G)(1 - \mu^G) \frac{\frac{\partial \xi}{\partial \mu^G}(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)}{\xi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G) \ln(\xi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G))}}{(1 - \mu^G) \ln(\xi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G))} \\ &= - \frac{\psi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)}{(1 - \mu^G) \ln(\xi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G))}. \quad \blacksquare \end{aligned}$$

The following two lemmas concern properties of $\psi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)$ that will be useful for the proof of Lemma 6. Lemma OA 2 says that $\psi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)$ is well-

defined and continuous in (γ, μ^G) , while Lemma OA 3 says that it has at most two zeros in $\mu^G \in (0, 1)$ when $\gamma = 1$. These lemmas, when combined with Lemma OA 1, imply that the magnitude of $\partial \tilde{K} / \partial \mu^G(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)$ becomes arbitrarily large for almost all $\mu^G \in (0, 1)$ as $\gamma \rightarrow 1$.

Lemma OA 2. $\psi : (0, 1] \times (0, 1) \times (0, 1) \times (0, 1) \rightarrow \mathbb{R}$ is well-defined and $\psi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)$ is continuous in (γ, μ^G) .

Proof of Lemma OA 2. Since $\xi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)$ is differentiable and only takes values in $(0, 1)$, $\psi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)$ is well-defined. Moreover, since $\xi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)$ is continuously differentiable for all $(\gamma, \mu^G) \in (0, 1) \times (0, 1)$, $\psi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)$ is continuous in (γ, μ^G) when $\gamma < 1$. All that remains is to check that $\psi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)$ is continuous in (γ, μ^G) at $\gamma = 1$.

First, note that $\psi(1, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)$ is continuous in μ^G . Thus, we need only check the limit in which γ approaches 1, but never equals 1. Note that

$$\begin{aligned} & \frac{\frac{\partial \xi}{\partial \mu^G}(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)}{\xi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G) \ln(\xi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G))} \\ &= \frac{\frac{\gamma(\tilde{\varepsilon}_{(a,a)} - \tilde{\varepsilon}_{(b,b)})(1-\gamma)}{(1-\gamma + \gamma(\tilde{\varepsilon}_{(a,a)}\mu^G + \tilde{\varepsilon}_{(b,b)}(1-\mu^G)))^2}}{\xi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G) \ln(\xi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G))} \quad (\text{OA 2}) \\ &= \left(\frac{\tilde{\varepsilon}_{(a,a)} - \tilde{\varepsilon}_{(b,b)}}{\tilde{\varepsilon}_{(a,a)}\mu^G + \tilde{\varepsilon}_{(b,b)}(1-\mu^G)} \right) \left(\frac{1 - \xi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)}{\ln(\xi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G))} \right). \end{aligned}$$

For γ close to 1,

$$\ln(\xi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)) = \xi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G) - 1 + O((\xi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G) - 1)^2).$$

Thus,

$$\lim_{\substack{(\gamma, \mu) \rightarrow (1, \mu^G) \\ \gamma \neq 1}} \frac{1 - \xi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)}{\ln(\xi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G))} = -1$$

for all $\mu^G \in (0, 1)$. This implies that $\psi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)$ is continuous in (γ, μ^G) at $\gamma = 1$. ■

Lemma OA 3. For any fixed ε , $\psi(1, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)$ has at most two zeros in $\mu^G \in (0, 1)$.

Proof of Lemma OA 3. When $\tilde{\varepsilon}_{(a,a)} \geq \tilde{\varepsilon}_{(b,b)}$, $\psi(1, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G) > 0$ for all $\mu^G \in (0, 1)$. When $\tilde{\varepsilon}_{(b,b)} > \tilde{\varepsilon}_{(a,a)}$, it suffices to show that

$$\frac{\ln(1 - \mu^G)(1 - \mu^G)}{\tilde{\varepsilon}_{(a,a)}\mu^G + \tilde{\varepsilon}_{(b,b)}(1 - \mu^G)}$$

is single-peaked in $\mu^G \in (0, 1)$. Note that

$$\frac{\partial}{\partial \mu^G} \left[\frac{\ln(1 - \mu^G)(1 - \mu^G)}{\tilde{\varepsilon}_{(a,a)}\mu^G + \tilde{\varepsilon}_{(b,b)}(1 - \mu^G)} \right] = \frac{(\tilde{\varepsilon}_{(b,b)} - \tilde{\varepsilon}_{(a,a)})\mu^G - \tilde{\varepsilon}_{(a,a)} \ln(1 - \mu^G) - \tilde{\varepsilon}_{(b,b)}}{(\tilde{\varepsilon}_{(a,a)}\mu^G + \tilde{\varepsilon}_{(b,b)}(1 - \mu^G))^2}.$$

The single-peakedness of $\ln(1 - \mu^G)(1 - \mu^G)/(\tilde{\varepsilon}_{(a,a)}\mu^G + \tilde{\varepsilon}_{(b,b)}(1 - \mu^G))$ follows from $(\tilde{\varepsilon}_{(b,b)} - \tilde{\varepsilon}_{(a,a)})\mu^G - \tilde{\varepsilon}_{(a,a)} \ln(1 - \mu^G) - \tilde{\varepsilon}_{(b,b)}$ being increasing in μ^G . \blacksquare

With these preliminaries established, we now present the proof of Lemma 6.

Proof of Lemma 6. Fix ε . Lemma OA 3 says $\psi(1, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)$ has at most two zeros for $\mu^G \in (0, 1)$. Because of this, there exists $\mu_1, \mu_2, \mu_3, \mu_4, \mu_5, \mu_6 \in (0, 1)$ satisfying $0 < \mu_1 < \mu_2 < \mu_3 < \mu_4 < \mu_5 < \mu_6 < 1$ such that

$$\min\{|\mu^G - \mu_1|, |\mu^G - \mu_2|, |\mu^G - \mu_3|, |\mu^G - \mu_4|, |\mu^G - \mu_5|, |\mu^G - \mu_6|\} < \Delta/2 \quad (\text{OA 3})$$

for all $\mu^G \in [0, 1]$, and $\psi(1, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)$ is non-zero on the intervals $[\mu_1, \mu_2]$, $[\mu_3, \mu_4]$, and $[\mu_5, \mu_6]$. Let $M = [\mu_1, \mu_2] \cup [\mu_3, \mu_4] \cup [\mu_5, \mu_6]$. Equation OA 3 says that the interval endpoints can be chosen so that M is no farther than $\Delta/2$ from any $\mu^G \in [0, 1]$, while the second condition implies that

$$|\psi(1, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)| > 0 \text{ for all } \mu^G \in M. \quad (\text{OA 4})$$

Lemma OA 2 says $\psi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)$ is continuous for $(\gamma, \mu^G) \in (0, 1] \times (0, 1)$. Hence, $\psi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)$ is uniformly continuous for $(\gamma, \mu^G) \in [\underline{\gamma}, 1] \times M$ for any

$\underline{\gamma} > 0$. Equation OA 4 then implies that there exists some $\lambda > 0$ and $\tilde{\gamma} \in (0, 1)$ such that $|\psi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)/(1 - \mu^G)| > \lambda$ for all $\gamma > \tilde{\gamma}$ and $\mu^G \in M$.

Define $\eta \in (0, 1)$ to be

$$\eta = \min \left\{ \frac{\mu_2 - \mu_1}{2}, \frac{\mu_4 - \mu_3}{2}, \frac{\mu_6 - \mu_5}{2}, \frac{\Delta}{2} \right\}.$$

Because $\lim_{\gamma \rightarrow 1} \min_{\mu^G \in [0,1]} \xi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G) = 1$, there exists $\gamma' \in (0, 1)$ such that $|\ln(\xi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G))| < \lambda\eta$ for all $\gamma > \gamma'$ and $\mu^G \in M$.

Moreover, $\lim_{\gamma \rightarrow 1} \min_{\mu^G \in [0,1]} \xi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G) = 1$ implies that there exists $\hat{\gamma} \in (0, 1)$ such that $\tilde{K}(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G) \geq 1$ for all $\gamma > \hat{\gamma}$ and $\mu^G \in M$.

Let $\bar{\gamma} = \max\{\tilde{\gamma}, \gamma', \hat{\gamma}\}$. Thus, $|\psi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)/((1 - \mu^G) \ln(\xi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)))| > 1/\eta$ and $\tilde{K}(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G) \geq 1$ for all $\gamma > \bar{\gamma}$ and $\mu^G \in M$. For the remainder of the proof, fix $\gamma \in (\bar{\gamma}, 1)$. We now show that, for a given $\mu^G \in M$, there exists some $\hat{\mu}^G \in M$ and non-negative integer \hat{K} such that $|\hat{\mu}^G - \mu^G| < \Delta/2$ and $\hat{\mu}^G = 1 - \xi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \hat{\mu}^G)^{\hat{K}}$. This, when combined with Equation OA 3, completes the proof.

Fix $\mu^G \in M$. Suppose for concreteness that $\mu^G \in [\mu_1, \mu_2]$. An identical argument handles the case when $\mu^G \in [\mu_3, \mu_4] \cup [\mu_5, \mu_6]$. By construction, η is weakly smaller than both $(\mu_2 - \mu_1)/2$ and $\Delta/2$. Therefore, there is some $\tilde{\mu}^G \in [\mu_1, \mu_2]$ such that $\eta \leq |\tilde{\mu}^G - \mu^G| \leq \Delta/2$. Because $|\psi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)/((1 - \mu^G) \ln(\xi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)))| > 1/\eta$ for all $\mu^G \in M$, it follows from Lemma OA 1 that $|\partial\xi/\partial\mu^G(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)| > 1/\eta$ for all $\mu^G \in M$. Hence, $|\tilde{K}(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \tilde{\mu}^G) - \tilde{K}(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)| > 1$. It thus follows that there exists some $\hat{\mu}^G$ between μ^G and $\tilde{\mu}^G$ and some non-negative integer \hat{K} between $\tilde{K}(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \mu^G)$ and $\tilde{K}(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \tilde{\mu}^G)$ such that $\tilde{K}(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \hat{\mu}^G) = \hat{K}$. Thus, $|\hat{\mu}^G - \mu^G| < \Delta/2$ and $\hat{\mu}^G = 1 - \xi(\gamma, \tilde{\varepsilon}_{(a,a)}, \tilde{\varepsilon}_{(b,b)}, \hat{\mu}^G)^{\hat{K}}$. \blacksquare

OA.2 Proof of Theorem 3

We show that action a can be limit-supported by the following class of tolerant grim trigger strategies. Let each player's score $k \in \mathbb{N}$ be the number of times their action was recorded as anything other than a or b . 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 against fellow good-standing players and play b against bad-standing players, while bad-standing players play c against good-standing players and play d against bad-standing players. We show that K can be appropriately tailored for parameters (γ, ε) so that a is limit-supported by strict equilibria in the limit $(\gamma, \varepsilon) \rightarrow (1, 0)$.

For every $k \in \{0, \dots, K - 1\}$, let μ_k be the share of players with score k and let $\mu^G = \sum_{k=0}^{K-1} \mu_k$ be the total share of players in good standing. The inflow into score 0 is simply $1 - \gamma$, the share of newborn players. The outflow from score 0 is the sum of $(1 - \gamma)\mu_0$, the share of players with score 0 who die in a given period, and $\gamma(n - 2)\varepsilon\mu_0$, the share of players who move from score 0 to score 1 in a given period. In a steady state, the inflow and outflow of score 0 must be equal, so we obtain $\mu_0 = (1 - \gamma)/(1 - \gamma(1 - (n - 2)\varepsilon))$. Similarly, for any $0 < k < K$, the inflow is $\gamma(n - 2)\varepsilon\mu_{k-1}$, the share of players who move from score $k - 1$ to score k in a given period. The outflow from score k is the sum of $(1 - \gamma)\mu_k$, the share of players with score k who die in a given period, and $\gamma(n - 2)\varepsilon\mu_k$, the share of players who move from score k to score $k + 1$ in a given period. Thus, in steady state, $\mu_k = \gamma(n - 2)\varepsilon\mu_{k-1}/(1 - \gamma(1 - (n - 2)\varepsilon))$. Combining these results gives $\mu_k = \nu(\gamma, \varepsilon)^k(1 - \nu(\gamma, \varepsilon))$ for every $k \in \{0, \dots, K - 1\}$, where $\nu : (0, 1) \times (0, 1) \rightarrow [0, 1]$ is the function given by

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

Summing $\mu_k = \nu(\gamma, \varepsilon)^k(1 - \nu(\gamma, \varepsilon))$ over $k \in \{0, \dots, K - 1\}$ gives

$$\mu^G = 1 - \nu(\gamma, \varepsilon)^K. \tag{OA 5}$$

It follows that there is a steady-state with parameters (γ, ε) , tolerance level K , and share of good-standing players $\mu^G \in [0, 1]$ precisely when Equation OA 5 holds.

Now we turn to the incentives of the good-standing players.¹ The expected payoff of a good-standing player with score k from playing a against a is $(1-\gamma)u(a, a) + \gamma(1-(n-2)\varepsilon)V_k + \gamma(n-2)\varepsilon V_{k+1}$, while their expected payoff from playing some action $a' \neq a, b$ is $(1-\gamma)u(a', a) + 2\gamma\varepsilon V_k + \gamma(1-2\varepsilon)V_{k+1}$.² Such a player will prefer to play a against a rather than playing action $a' \neq b$ if $(1-n\varepsilon)\gamma(V_k - V_{k+1})/(1-\gamma) > u(a', a) - u(a, a)$. Similarly, the incentives of a good-standing player with score k to play b against c are satisfied if $(1-n\varepsilon)\gamma(V_k - V_{k+1})/(1-\gamma) > u(a', c) - u(b, c)$ for all $a' \neq a, b$. Thus, it follows that the incentives of good-standing players both to play a against a and play b against c are satisfied if $(1-n\varepsilon)\gamma(V_k - V_{k+1})/(1-\gamma)$ is sufficiently high.

For $k < K$, the relationship between the value function at score k and the value function at score $k+1$ is $V_k = (1-\gamma)(\mu^G u(a, a) + (1-\mu^G)u(b, c)) + \gamma(1-(n-2)\varepsilon)V_k + \gamma(n-2)\varepsilon V_{k+1}$. Combining this with the fact that the bad-standing value functions are $V_k = \mu^G u(c, b) + (1-\mu^G)u(d, d)$ for $k \geq K$, it follows that

$$\begin{aligned} V_k = & (1 - \nu(\gamma, \varepsilon)^{K-k})(\mu^G u(a, a) + (1 - \mu^G)u(b, c)) \\ & + \nu(\gamma, \varepsilon)^{K-k}(\mu^G u(c, b) + (1 - \mu^G)u(d, d)) \end{aligned} \quad (\text{OA } 6)$$

for $k < K$. An implication of Equation OA 6 is that $V_0 - V_1 \leq V_k - V_{k+1}$ for all $k < K$, so the incentives of good-standing players with score 0 will be the most difficult to satisfy. Moreover, combining Equations OA 5 and OA 6 gives

$$\begin{aligned} & (1 - n\varepsilon) \frac{\gamma}{1 - \gamma} (V_0 - V_1) \\ = & \frac{(1 - n\varepsilon)(1 - \mu^G)}{(n - 2)\varepsilon} (\mu^G (u(a, a) - u(c, b)) + (1 - \mu^G)(u(b, c) - u(d, d))). \end{aligned} \quad (\text{OA } 7)$$

Fix $\eta \in (0, 1)$. For arbitrary (γ, ε) , consider the tolerant grim trigger strategy with

¹Since c is a strict best-response to b and (d, d) is a strict static equilibrium, the incentives of the bad-standing players are always satisfied.

²Because $u(a, a) > u(b, a)$, a good-standing player never benefits from playing b rather than a against an opponent playing a .

tolerance $K(\gamma, \varepsilon)$ set by

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

so that $K(\gamma, \varepsilon)$ is the smallest integer greater than or equal to $\ln(\eta)/\ln(\nu(\gamma, \varepsilon))$. Note that with this tolerance the corresponding share of good-standing players is $\mu^G(\gamma, \varepsilon) = 1 - \nu(\gamma, \varepsilon)^{K(\gamma, \varepsilon)}$. 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 tolerant grim trigger strategy is an equilibrium whenever $K(\gamma, \varepsilon) > 1$.

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

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

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 tolerant grim trigger strategy is an equilibrium whenever $K(\gamma, \varepsilon) = 1$.

OA.3 Proof of Lemma 16

Lemma 16. *There is a $P_K S_L D_\infty$ equilibrium with share of cooperators μ^C , share of reciprocators μ^P , and share of supercooperators μ^S if and only if the following conditions hold:*

1. *Feasibility:*

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

2. *Incentives:*

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

We first derive the feasibility conditions and then derive the incentive conditions. The feasibility conditions of Lemma 16 follow from the following lemma.

Lemma OA 4. *In a $P_K S_L D_\infty$ steady state with total share of cooperators μ^C ,*

$$\mu_k = \begin{cases} \beta(\gamma, \varepsilon, \mu^C)^k (1 - \beta(\gamma, \varepsilon, \mu^C)) & \text{if } 0 \leq k \leq K - 1 \\ \alpha(\gamma, \varepsilon_C)^{k-K} \beta(\gamma, \varepsilon, \mu^C)^K (1 - \alpha(\gamma, \varepsilon_C)) & \text{if } K \leq k \leq K + L - 1 \end{cases}.$$

To see why Lemma OA 4 implies the feasibility conditions of Lemma 16, note that

$$\begin{aligned}\mu^P &= \sum_{k=0}^{K-1} \beta(\gamma, \varepsilon, \mu^C)^k (1 - \beta(\gamma, \varepsilon, \mu^C)) = 1 - \beta(\gamma, \varepsilon, \mu^C)^K, \\ \mu^S &= \sum_{k=K}^{K+L-1} \alpha(\gamma, \varepsilon_C)^{k-K} \beta(\gamma, \varepsilon, \mu^C)^K (1 - \alpha(\gamma, \varepsilon_C)) = (1 - \alpha(\gamma, \varepsilon_C)^L) \beta(\gamma, \varepsilon, \mu^C)^K,\end{aligned}$$

which also gives $\mu^C = \mu^P + \mu^S = 1 - \alpha(\gamma, \varepsilon_C)^L \beta(\gamma, \varepsilon, \mu^C)^K$.

Proof of Lemma OA 4. The inflow into record 0 is $1 - \gamma$, while the outflow from record 0 is $(1 - \gamma(\varepsilon_D + (1 - \varepsilon_C - \varepsilon_D)\mu^C))\mu_0$. Setting these equal gives

$$\mu_0 = \frac{1 - \gamma}{1 - \gamma(\varepsilon_D + (1 - \varepsilon_C - \varepsilon_D)\mu^C)} = 1 - \beta(\gamma, \varepsilon, \mu^C).$$

Additionally, for every $0 < k < K$, both record k and record $k - 1$ are preciprocators. Thus, the inflow into record k is $\gamma(1 - \varepsilon_D - (1 - \varepsilon_C - \varepsilon_D)\mu^C)\mu_{k-1}$, while the outflow from record k is $(1 - \gamma(\varepsilon_D + (1 - \varepsilon)\mu^C))\mu_k$. Setting these equal gives

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

Combining this with $\mu_0 = 1 - \beta(\gamma, \varepsilon, \mu^C)$ gives $\mu_k = \beta(\gamma, \varepsilon, \mu^C)^k(1 - \beta(\gamma, \varepsilon, \mu^C))$ for $0 \leq k \leq K - 1$. Note that this is sufficient to establish Lemma OA 4 for the case where there are no supercooperator records ($L = 0$). For the rest of the proof, we assume that $L > 0$.

Since record $K - 1$ is a preciprocator and record K is a supercooperator, the inflow into record K is $\gamma(1 - \varepsilon_D - (1 - \varepsilon_C - \varepsilon_D)\mu^C)\mu_{K-1}$, while the outflow is $(1 - \gamma(1 - \varepsilon_C))\mu_K$. Setting these equal and using the fact that $\mu_{K-1} = \beta(\gamma, \varepsilon, \mu^C)^{K-1}(1 - \beta(\gamma, \varepsilon, \mu^C))$, we have

$$\begin{aligned} \mu_K &= \frac{\gamma(1 - \varepsilon_D - (1 - \varepsilon_C - \varepsilon_D)\mu^C)}{1 - \gamma(1 - \varepsilon)}\beta(\gamma, \varepsilon, \mu^C)^{K-1}(1 - \beta(\gamma, \varepsilon, \mu^C)) \\ &= \beta(\gamma, \varepsilon, \mu^C)^K \frac{1 - \gamma}{1 - \gamma(1 - \varepsilon_C)} \\ &= \beta(\gamma, \varepsilon, \mu^C)^K(1 - \alpha(\gamma, \varepsilon_C)). \end{aligned}$$

Additionally, for every $K < k < K + L$, both record k and record $k - 1$ are supercooperators. Thus, the inflow into record k is $\gamma\varepsilon_C\mu_{k-1}$, while the outflow from record k is $(1 - \gamma(1 - \varepsilon_C))\mu_k$. Setting these equal gives

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

Combining this with $\mu_K = \beta(\gamma, \varepsilon, \mu^C)^K(1 - \alpha(\gamma, \varepsilon_C))$ gives $\mu_k = \alpha(\gamma, \varepsilon_C)^{k-K}\beta(\gamma, \varepsilon, \mu^C)^K(1 - \alpha(\gamma, \varepsilon_C))$ for all $K \leq k \leq K + L - 1$. ■

Now we establish the incentives conditions in Lemma 16. We first handle the incentives of the record 0 preciprocator. Since V_0 equals the average payoff in the

population in every period, we have

$$V_0 = \mu^P \mu^C + \mu^S(\mu^C - \mu^D l) + \mu^D \mu^S(1 + g).$$

Since the flow payoff to a reciprocator is μ^C , Lemma 8 along with the fact that $p_k^D = 1 - \varepsilon_D - (1 - \varepsilon_C - \varepsilon_D)\mu^C$ for any reciprocator record implies that a record 0 reciprocator plays C against C iff

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

Since

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

it follows that the $(C|C)_0$ constraint is equivalent to

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

Now, we handle the incentives of a record K supercooperator. Since V_K equals the average payoff experienced by players in the population with a record $k \geq K$, we have

$$V_K = \frac{1}{1 - \mu^P} (\mu^S(\mu^C - \mu^D l) + \mu^D \mu^S(1 + g)).$$

Since the flow payoff to a supercooperator is $\mu^C - \mu^D l$, Lemma 8 along with the fact that $p_k^D = \varepsilon_C$ for any supercooperator record implies that a record K supercooperator plays C against D iff

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

Note that this is equivalent to

$$(1 - \varepsilon_C - \varepsilon_D)(1 - \mu^C)(\mu^P + \mu^S(l - g)) + (1 - \varepsilon_C - \varepsilon_D)(1 - \mu^C)l(1 - \mu^P) > \varepsilon_C l(1 - \mu^P).$$

Further manipulating this gives

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

from which the $(C|D)_K$ in Lemma 16 immediately follows.

Finally, we handle the incentives of a record $K - 1$ preciprocator. Note that

$$V_{K-1} = (1 - \gamma)\mu^C + \gamma((1 - \varepsilon_C)\mu^C + \varepsilon_D(1 - \mu^C))V_{K-1} + \gamma(\varepsilon_C\mu^C + (1 - \varepsilon_D)(1 - \mu^C))V_K,$$

so

$$\frac{\gamma(1 - \varepsilon_C - \varepsilon_D)}{1 - \gamma}(V_{K-1} - V_K) = \frac{\gamma(1 - \varepsilon_C - \varepsilon_D)}{1 - \gamma(\varepsilon_D + (1 - \varepsilon_C - \varepsilon_D)\mu^C)}(\mu^C - V_K).$$

Since the value function for a record K player is

$$V_K = \frac{1}{1 - \mu^P}(\mu^S(\mu^C - \mu^D l) + \mu^D \mu^S(1 + g)),$$

we find

$$\begin{aligned} & \frac{\gamma(1 - \varepsilon_C - \varepsilon_D)}{1 - \gamma(\varepsilon_D + (1 - \varepsilon_C - \varepsilon_D)\mu^C)}(\mu^C - V_K) \\ &= \frac{\gamma(1 - \varepsilon_C - \varepsilon_D)}{1 - \gamma(\varepsilon_D + (1 - \varepsilon_C - \varepsilon_D)\mu^C)} \left(\mu^C - \frac{1}{1 - \mu^P}(\mu^S(\mu^C - \mu^D l) + \mu^D \mu^S(1 + g)) \right) \\ &= \frac{\gamma(1 - \varepsilon_C - \varepsilon_D)(1 - \mu^C)}{1 - \gamma(\varepsilon_D + (1 - \varepsilon_C - \varepsilon_D)\mu^C)} \frac{1}{1 - \mu^P}(\mu^P + \mu^S(l - g)). \end{aligned}$$

Thus, a record $K - 1$ preciprocator plays D against D iff

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

which implies the form of the $(D|D)_{K-1}$ constraint in Lemma 16.

OA.4 Proof of Lemma 18

Lemma 18. *If $l > \max\{g + g^2, f(g)\}$, there is some $\bar{\varepsilon} > 0$ and $\mu^P \in (g/(1+g), 1 - g/l]$ such that $\liminf_{\gamma \rightarrow 1} \bar{\mu}_{PSD}^C(\gamma, \varepsilon) \geq h(\varepsilon, \mu^P)$ for $\tilde{\varepsilon}_C, \tilde{\varepsilon}_D < \bar{\varepsilon}$.*

Lemma 18 is a consequence of the following two lemmas. Let $\kappa : (g/(1+g), 1 - g/l] \rightarrow \mathbb{R}$ be the function given by

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

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

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

Lemma OA 5. *Fix $\mu^P \in (g/(1+g), 1 - g/l]$. If $|1 + \kappa(\mu^P)| > \iota(\mu^P)$, then there exists some $\bar{\varepsilon} > 0$ such that $\liminf_{\gamma \rightarrow 1} \bar{\mu}_{PSD}^C(\gamma, \varepsilon) \geq h(\varepsilon, \mu^P)$ for $\varepsilon_C, \varepsilon_D < \bar{\varepsilon}$.*

Lemma OA 6. *If $l > \max\{g + g^2, f(g)\}$, some $\mu^P \in (g/(1+g), 1 - g/l]$ satisfies $|1 + \kappa(\mu^P)| > \iota(\mu^P)$.*

OA.4.1 presents the proof of Lemma OA 5. It uses the inverse function theorem and other tools of differential calculus to show that, when $|1 + \kappa(\mu^P)| > \iota(\mu^P)$, for sufficiently small ε_C and ε_D , any neighborhood of $(h(\varepsilon, \mu^P), \mu^P)$ can be approached by feasible profiles for sufficiently high γ . This proof is fairly long, as it proceeds through a series of subsidiary lemmas that concludes with Lemma OA 14. The proof of Lemma OA 6 is in OA.4.2.

OA.4.1 Proof of Lemma OA 5

Let $\rho : [0, 1] \times (0, 1) \times (0, 1) \times [0, 1] \rightarrow [0, 1]$ be the function given by

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

Equation 13 can be equivalently written as

$$\rho(1, \varepsilon, \mu^C) ((l - g)\mu^C + (1 + g - l)\mu^P) + l\mu^P = l.$$

Setting $\mu^C = h(\varepsilon, \mu^P)$ in the above equation and solving for $\rho(1, \varepsilon, h(\varepsilon, \mu^P))$ gives

$$\rho(1, \varepsilon, h(\varepsilon, \mu^P)) = \frac{l(1 - \mu^P)}{(l - g)h(\varepsilon, \mu^P) + (1 + g - l)\mu^P}$$

for all $\varepsilon_C, \varepsilon_D$ such that $h(\varepsilon, \mu^P)$ is well-defined. Since $\lim_{\varepsilon \rightarrow 0} h(\varepsilon, \mu^P) = 1$, an immediate corollary follows.

Corollary OA 1. *For every $\mu^P \in (g/(1 + g), 1 - g/l]$,*

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

Define the function $I : [0, 1] \times (0, 1) \times (0, 1) \times [0, 1] \times [0, 1] \rightarrow \mathbb{R}$ by

$$I(\gamma, \varepsilon, \mu^C, \mu^P) = \rho(\gamma, \varepsilon, \mu^C) ((l - g)\mu^C + (1 + g - l)\mu^P) + l\mu^P.$$

The $(D|D)_{K-1}$ constraint is equivalent to $I(\gamma, \varepsilon, \mu^C, \mu^P) < l$, and the $(C|D)_K$ constraint is equivalent to $I(1, \varepsilon, \mu^C, \mu^P) > l$. The $(C|C)_0$ constraint holds whenever the $(C|D)_K$ constraint holds and $\mu^P \leq 1 - g/l$, which is true for the profiles we consider.

Lemma OA 7. *Fix $\mu^P \in (g/(1 + g), 1 - g/l]$. There exists $\bar{\varepsilon} > 0$ such that*

$$\frac{\partial I}{\partial \mu^C}(1, \varepsilon, h(\varepsilon, \mu^P), \mu^P) < 0 < \frac{\partial I}{\partial \mu^P}(1, \varepsilon, h(\varepsilon, \mu^P), \mu^P)$$

for all $\varepsilon_C, \varepsilon_D < \bar{\varepsilon}$.

Proof of Lemma OA 7. Note that

$$\begin{aligned}\frac{\partial I}{\partial \mu^P}(1, \varepsilon, h(\varepsilon, \mu^P), \mu^P) &= \rho(1, \varepsilon, h(\varepsilon, \mu^P))(1 + g - l) + l \\ &> \rho(1, \varepsilon, h(\varepsilon, \mu^P))(1 + g) > 0,\end{aligned}$$

since $0 < \rho(1, \varepsilon, h(\varepsilon, \mu^P)) < 1$.

Moreover,

$$\begin{aligned}&\frac{\partial I}{\partial \mu^C}(1, \varepsilon, h(\varepsilon, \mu^P), \mu^P) \\ &= - \left(\frac{1}{1 - h(\varepsilon, \mu^P)} \right) \left(\frac{\varepsilon_C}{1 - \varepsilon_D - (1 - \varepsilon_C - \varepsilon_D)h(\varepsilon, \mu^P)} \right) \\ &\quad \rho(1, \varepsilon, h(\varepsilon, \mu^P))((1 + g - l)\mu^P + (l - g)h(\varepsilon, \mu^P)) \\ &\quad + \rho(1, \varepsilon, h(\varepsilon, \mu^P))(l - g) \\ &= - \left(\frac{1}{1 - h(\varepsilon, \mu^P)} \right) \left(\frac{\varepsilon_C}{1 - \varepsilon_D - (1 - \varepsilon_C - \varepsilon_D)h(\varepsilon, \mu^P)} \right) l(1 - \mu^P) \\ &\quad + \rho(1, \varepsilon, h(\varepsilon, \mu^P))(l - g).\end{aligned}$$

Combining $\lim_{\varepsilon \rightarrow 0} h(\varepsilon, \mu^P) = 1$ and Equation 14 gives

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

It follows that

$$\lim_{\varepsilon \rightarrow 0} \frac{\partial I}{\partial \mu^C}(1, \varepsilon, h(\varepsilon, \mu^P), \mu^P) = -\infty.$$

Thus, there exists some $\bar{\varepsilon} > 0$ such that

$$\frac{\partial I}{\partial \mu^C}(1, \varepsilon, h(\varepsilon, \mu^P), \mu^P) < 0$$

for all $\varepsilon_C, \varepsilon_D < \bar{\varepsilon}$. ■

Let $\tilde{K} : (0, 1) \times (0, 1) \times (0, 1) \times (0, 1) \times (0, 1) \rightarrow \mathbb{R}$ be the function given by

$$\tilde{K}(\gamma, \varepsilon, \mu^C, \mu^P) = \frac{\ln(1 - \mu^P)}{\ln(\beta(\gamma, \varepsilon, \mu^C))}, \quad (\text{OA } 8)$$

and $\tilde{L} : (0, 1) \times (0, 1) \times (0, 1) \times (0, 1) \times (0, 1) \rightarrow \mathbb{R}_+$ be the function given by

$$\tilde{L}(\gamma, \varepsilon, \mu^C, \mu^P) = \frac{\ln(1 - \mu^C) - \ln(1 - \mu^P)}{\ln(\alpha(\gamma, \varepsilon_C))}. \quad (\text{OA } 9)$$

Note that $\tilde{L}(\gamma, \varepsilon, \mu^C, \mu^P) \geq 0$ whenever $\mu^C \geq \mu^P$, which is the case of interest. By construction, $\tilde{K}(\gamma, \varepsilon, \mu^C, \mu^P)$ and $\tilde{L}(\gamma, \varepsilon, \mu^C, \mu^P)$ are the unique $(K, L) \in \mathbb{R}_+ \times \mathbb{R}_+$ such that the feasibility constraints in Lemma 16 are satisfied.

Differentiating Equations OA 8 and OA 9 gives the following result.

Lemma OA 8. $\tilde{K}(\gamma, \varepsilon, \mu^C, \mu^P)$ and $\tilde{L}(\gamma, \varepsilon, \mu^C, \mu^P)$ are differentiable in $(\mu^C, \mu^P) \in (0, 1) \times (0, 1)$ with partial derivatives

$$\begin{aligned} \frac{\partial \tilde{K}}{\partial \mu^C}(\gamma, \varepsilon, \mu^C, \mu^P) &= -\frac{\ln(1 - \mu^P) \frac{\partial \beta}{\partial \mu}(\gamma, \varepsilon, \mu^C)}{\ln(\beta(\gamma, \varepsilon, \mu^C))^2 \beta(\gamma, \varepsilon, \mu^C)}, \\ \frac{\partial \tilde{L}}{\partial \mu^C}(\gamma, \varepsilon, \mu^C, \mu^P) &= -\frac{1}{(1 - \mu^C) \ln(\alpha(\gamma, \varepsilon_C))}, \\ \frac{\partial \tilde{K}}{\partial \mu^P}(\gamma, \varepsilon, \mu^C, \mu^P) &= -\frac{1}{(1 - \mu^P) \ln(\beta(\gamma, \varepsilon, \mu^C))}, \\ \frac{\partial \tilde{L}}{\partial \mu^P}(\gamma, \varepsilon, \mu^C, \mu^P) &= \frac{1}{(1 - \mu^P) \ln(\alpha(\gamma, \varepsilon_C))}. \end{aligned}$$

Let $J(\gamma, \varepsilon, \mu^C, \mu^P)$ be the Jacobian matrix comprising the various partial derivatives of \tilde{K} and \tilde{L} . That is,

$$\begin{aligned} J(\gamma, \varepsilon, \mu^C, \mu^P) &= \begin{bmatrix} \frac{\partial \tilde{K}}{\partial \mu^C}(\gamma, \varepsilon, \mu^C, \mu^P) & \frac{\partial \tilde{K}}{\partial \mu^P}(\gamma, \varepsilon, \mu^C, \mu^P) \\ \frac{\partial \tilde{L}}{\partial \mu^C}(\gamma, \varepsilon, \mu^C, \mu^P) & \frac{\partial \tilde{L}}{\partial \mu^P}(\gamma, \varepsilon, \mu^C, \mu^P) \end{bmatrix} \\ &= \begin{bmatrix} -\frac{\ln(1 - \mu^P) \frac{\partial \beta}{\partial \mu}(\gamma, \varepsilon, \mu^C)}{\ln(\beta(\gamma, \varepsilon, \mu^C))^2 \beta(\gamma, \varepsilon, \mu^C)} & -\frac{1}{(1 - \mu^P) \ln(\beta(\gamma, \varepsilon, \mu^C))} \\ -\frac{1}{(1 - \mu^C) \ln(\alpha(\gamma, \varepsilon_C))} & \frac{1}{(1 - \mu^P) \ln(\alpha(\gamma, \varepsilon_C))} \end{bmatrix}. \end{aligned}$$

Let $\zeta : [0, 1] \times (0, 1) \times (0, 1) \times (0, 1) \times (0, 1) \rightarrow \mathbb{R}$ be the function given by

$$\zeta(\gamma, \varepsilon, \mu^C, \mu^P) = \begin{cases} \ln(1 - \mu^P) \frac{(1 - \mu^C) \frac{\partial \beta}{\partial \mu}(\gamma, \varepsilon, \mu^C)}{\beta(\gamma, \varepsilon, \mu^C) \ln(\beta(\gamma, \varepsilon, \mu^C))} & \text{if } \gamma < 1 \\ \ln(1 - \mu^P) \rho(1, \varepsilon, \mu^C) & \text{if } \gamma = 1 \end{cases}.$$

The following lemma comes from direct calculation.

Lemma OA 9.

1. The determinant of $J(\gamma, \varepsilon, \mu^C, \mu^P)$ is

$$\det(J(\gamma, \varepsilon, \mu^C, \mu^P)) = -\frac{1 + \zeta(\gamma, \varepsilon, \mu^C, \mu^P)}{(1 - \mu^C)(1 - \mu^P) \ln(\alpha(\gamma, \varepsilon_C)) \ln(\beta(\gamma, \varepsilon, \mu^C))}.$$

2. When $J(\gamma, \varepsilon, \mu^C, \mu^P)$ is invertible, its inverse is

$$J(\gamma, \varepsilon, \mu^C, \mu^P)^{-1} = \begin{bmatrix} -\frac{(1 - \mu^C) \ln(\beta(\gamma, \varepsilon, \mu^C))}{1 + \zeta(\gamma, \varepsilon, \mu^C, \mu^P)} & -\frac{(1 - \mu^C) \ln(\alpha(\gamma, \varepsilon_C))}{1 + \zeta(\gamma, \varepsilon, \mu^C, \mu^P)} \\ -\frac{(1 - \mu^P) \ln(\beta(\gamma, \varepsilon, \mu^C))}{1 + \zeta(\gamma, \varepsilon, \mu^C, \mu^P)} & \frac{\zeta(\gamma, \varepsilon, \mu^C, \mu^P)(1 - \mu^P) \ln(\alpha(\gamma, \varepsilon_C))}{1 + \zeta(\gamma, \varepsilon, \mu^C, \mu^P)} \end{bmatrix}.$$

We establish the continuity of $\zeta(\gamma, \varepsilon, \mu^P, \mu^C)$.

Lemma OA 10. For all $\varepsilon \in (0, 1)$, $\zeta(\gamma, \varepsilon, \mu^C, \mu^P)$ is continuous in (γ, μ^C, μ^P) .

Proof of Lemma OA 10. Clearly, $\zeta(\gamma, \varepsilon, \mu^C, \mu^P)$ is continuous whenever $\gamma < 1$. What remains is to show that it is continuous when $\gamma = 1$. Note that $\ln(1 - \mu^P) \rho(1, \varepsilon, \mu^C)$ is continuous in (μ^C, μ^P) . Thus, we need only check the limit in which γ approaches 1, but never equals 1. Recall that

$$\begin{aligned} & \frac{\frac{\partial \beta}{\partial \mu}(\gamma, \varepsilon, \mu^C)}{\beta(\gamma, \varepsilon, \mu^C) \ln(\beta(\gamma, \varepsilon, \mu^C))} \\ &= -\frac{\frac{\gamma(1 - \varepsilon_C - \varepsilon_D)(1 - \gamma)}{(1 - \gamma(\varepsilon_D + (1 - \varepsilon_C - \varepsilon_D)\mu^C))^2}}{\beta(\gamma, \varepsilon, \mu^C) \ln(\beta(\gamma, \varepsilon, \mu^C))} \\ &= -\left(\frac{\gamma(1 - \varepsilon_C - \varepsilon_D)}{\beta(\gamma, \varepsilon, \mu^C)(1 - \gamma(\varepsilon_D + (1 - \varepsilon_C - \varepsilon_D)\mu))} \right) \left(\frac{1 - \beta(\gamma, \varepsilon, \mu^C)}{\ln(\beta(\gamma, \varepsilon, \mu^C))} \right). \end{aligned}$$

It is clear that

$$\lim_{\substack{(\gamma, \mu) \rightarrow (1, \mu^C) \\ \gamma \neq 1}} \frac{\gamma(1 - \varepsilon_C - \varepsilon_D)}{\beta(\gamma, \varepsilon, \mu)(1 - \gamma(\varepsilon_D + (1 - \varepsilon_C - \varepsilon_D)\mu))} = \frac{1 - \varepsilon_C - \varepsilon_D}{1 - \varepsilon_D - (1 - \varepsilon_C - \varepsilon_D)\mu^C}$$

for all $\mu^C \in (0, 1)$. For γ close to 1,

$$\ln(\beta(\gamma, \varepsilon, \mu^C)) = \beta(\gamma, \varepsilon, \mu^C) - 1 + O((\beta(\gamma, \varepsilon, \mu^C) - 1)^2).$$

Thus,

$$\lim_{\substack{(\gamma, \mu) \rightarrow (1, \mu^C) \\ \gamma \neq 1}} \frac{1 - \beta(\gamma, \varepsilon, \mu)}{\ln(\beta(\gamma, \varepsilon, \mu))} = -1$$

for all $\mu^C \in (0, 1)$. Combining these results implies

$$\lim_{\substack{(\gamma, \mu) \rightarrow (1, \mu^C) \\ \gamma \neq 1}} \frac{(1 - \mu) \frac{\partial \beta}{\partial \mu}(\gamma, \varepsilon, \mu)}{\beta(\gamma, \varepsilon, \mu) \ln(\beta(\gamma, \varepsilon, \mu))} = \rho(1, \varepsilon, \mu^C)$$

for all $\mu^C \in (0, 1)$. Hence, $\zeta(\gamma, \varepsilon, \mu^C, \mu^P)$ is continuous. ■

The following lemma concerns the extent to which, for small $\varepsilon_C, \varepsilon_D$ and fixed $\hat{\mu}^P \in (g/(1+g), 1-g/l]$, profiles (μ^C, μ^P) near $(h(\varepsilon, \hat{\mu}^P), \hat{\mu}^P)$ are close to feasible profiles. It combines Lemmas OA 9 and OA 10 with the inverse function theorem to obtain a bound on how far such (μ^C, μ^P) are from feasible profiles when the corresponding value of \tilde{L} is an integer. Moreover, the size of this bound is related to the magnitude of $1 + \zeta(1, \varepsilon, h(\varepsilon, \hat{\mu}^P), \hat{\mu}^P)$, which is close to $|1 + \kappa(\hat{\mu}^P)|$ for small $\varepsilon_C, \varepsilon_D$.

Lemma OA 11. *Fix $\hat{\mu}^P \in (g/(1+g), 1-g/l]$ and $\eta > 0$. If $|1 + \kappa(\hat{\mu}^P)| > \lambda$ for some $\lambda > 0$, there exists $\bar{\varepsilon} > 0$ such that, for all $\varepsilon_C, \varepsilon_D < \bar{\varepsilon}$, there is $\bar{\gamma} < 1$ and an open neighborhood of $(h(\varepsilon, \hat{\mu}^P), \hat{\mu}^P)$, M , where, for all $\gamma > \bar{\gamma}$ and $(\mu^C, \mu^P) \in M$, there are*

feasible $\tilde{\mu}^C$ and $\tilde{\mu}^P$ satisfying

$$\begin{aligned} 0 &\leq \tilde{\mu}^C - \mu^C < -\frac{1+\eta}{\lambda}(1 - h(\varepsilon, \hat{\mu}^P)) \ln(\beta(\gamma, \varepsilon, h(\varepsilon, \hat{\mu}^P))), \\ 0 &\leq \tilde{\mu}^P - \mu^P < -\frac{1+\eta}{\lambda}(1 - \hat{\mu}^P) \ln(\beta(\gamma, \varepsilon, h(\varepsilon, \hat{\mu}^P))), \end{aligned}$$

whenever $\tilde{L}(\gamma, \varepsilon, \mu^C, \mu^P)$ is an integer.

Proof of Lemma OA 11. We handle the case where $1 + \kappa(\hat{\mu}^P) > \lambda > 0$. The case where $1 + \kappa(\hat{\mu}^P) < -\lambda < 0$ can be handled analogously.

Note that

$$1 + \zeta(1, \varepsilon, h(\varepsilon, \hat{\mu}^P), \hat{\mu}^P) = 1 + \ln(1 - \hat{\mu}^P)\rho(1, \varepsilon, h(\varepsilon, \hat{\mu}^P)).$$

Moreover,

$$\lim_{\varepsilon \rightarrow 0} \ln(1 - \hat{\mu}^P)\rho(1, \varepsilon, h(\varepsilon, \hat{\mu}^P)) = \kappa(\hat{\mu}^P)$$

by Lemma OA 1. Thus, when $1 + \kappa(\hat{\mu}^P) > \lambda$, there exists some $\bar{\varepsilon} > 0$ such that, for all $\varepsilon_C, \varepsilon_D < \bar{\varepsilon}$, there exists $\bar{\gamma}_1 < 1$ and an open neighborhood of $(h(\varepsilon, \hat{\mu}^P), \hat{\mu}^P)$, M_1 , such that

$$1 + \zeta(\gamma, \varepsilon, \mu^C, \mu^P) < -\lambda$$

for all $\gamma > \bar{\gamma}_1$ and $(\mu^C, \mu^P) \in M_1$. By Lemma OA 9, $J(\gamma, \varepsilon, \mu^C, \mu^P)$ is invertible for all such points. Thus, for a given $\varepsilon_C, \varepsilon_D < \bar{\varepsilon}$ and $\gamma > \bar{\gamma}_1$, the inverse function theorem implies the existence of differentiable functions of (K, L) , $\tilde{\mu}^C$ and $\tilde{\mu}^P$, that constitute a local inverse of \tilde{K} and \tilde{L} for $(\mu^C, \mu^P) \in M_1$. Additionally, the partial derivatives of

these functions are given by J^{-1} , so that

$$\begin{aligned}\frac{\partial \tilde{\mu}^C}{\partial K}(\gamma, \varepsilon, K, L) &= -\frac{(1 - \tilde{\mu}^C(\gamma, \varepsilon, K, L)) \ln(\beta(\gamma, \varepsilon, \tilde{\mu}^C(\gamma, \varepsilon, K, L)))}{1 + \zeta(\gamma, \varepsilon, \tilde{\mu}^C(\gamma, \varepsilon, K, L), \tilde{\mu}^P(\gamma, \varepsilon, K, L))}, \\ \frac{\partial \tilde{\mu}^P}{\partial K}(\gamma, \varepsilon, K, L) &= -\frac{(1 - \tilde{\mu}^P(\gamma, \varepsilon, K, L)) \ln(\beta(\gamma, \varepsilon, \tilde{\mu}^C(\gamma, \varepsilon, K, L)))}{1 + \zeta(\gamma, \varepsilon, \tilde{\mu}^C(\gamma, \varepsilon, K, L), \tilde{\mu}^P(\gamma, \varepsilon, K, L))}, \\ \frac{\partial \tilde{\mu}^C}{\partial L}(\gamma, \varepsilon, K, L) &= -\frac{(1 - \tilde{\mu}^C(\gamma, \varepsilon, K, L)) \ln(\alpha(\gamma, \varepsilon_C))}{1 + \zeta(\gamma, \varepsilon, \tilde{\mu}^C(\gamma, \varepsilon, K, L), \tilde{\mu}^P(\gamma, \varepsilon, K, L))}, \\ \frac{\partial \tilde{\mu}^P}{\partial L}(\gamma, \varepsilon, K, L) &= \left(\frac{\zeta(\gamma, \varepsilon, \tilde{\mu}^C(\gamma, \varepsilon, K, L), \tilde{\mu}^P(\gamma, \varepsilon, K, L))}{1 + \zeta(\gamma, \varepsilon, \tilde{\mu}^C(\gamma, \varepsilon, K, L), \tilde{\mu}^P(\gamma, \varepsilon, K, L))} \right) \\ &\quad (1 - \tilde{\mu}^P(\gamma, \varepsilon, K, L)) \ln(\alpha(\gamma, \varepsilon_C)),\end{aligned}$$

for any (K, L) that equals $(\tilde{K}(\gamma, \varepsilon, \mu^C, \mu^P), \tilde{L}(\gamma, \varepsilon, \mu^C, \mu^P))$ for some $(\mu^C, \mu^P) \in M_1$.

There is a neighborhood of $(h(\varepsilon, \hat{\mu}^P), \hat{\mu}^P)$, M_2 , such that

$$1 - \mu^C < \sqrt{1 + \eta}(1 - h(\varepsilon, \hat{\mu}^P))$$

and

$$1 - \mu^P < \sqrt{1 + \eta}(1 - \hat{\mu}^P)$$

for all $(\mu^C, \mu^P) \in M_2$. Moreover, because $\beta(\gamma, \varepsilon, \mu^C)$ is decreasing in μ^C and

$$\lim_{\gamma \rightarrow 1} \frac{\ln(\beta(\gamma, \varepsilon, \mu_1^C))}{\ln(\beta(\gamma, \varepsilon, \mu_2^C))} = \frac{1 - \varepsilon_D - (1 - \varepsilon_C - \varepsilon_D)\mu_2^C}{1 - \varepsilon_D - (1 - \varepsilon_C - \varepsilon_D)\mu_1^C}$$

for all $(\gamma, \varepsilon) \in (0, 1) \times (0, 1)$ and $\mu_1^C, \mu_2^C \in [0, 1]$, we can take the neighborhood M_2 to be small enough so that

$$\ln(\beta(\gamma, \varepsilon, \mu^C)) > \sqrt{1 + \eta} \ln(\beta(\gamma, \varepsilon, h(\varepsilon, \hat{\mu}^P)))$$

for all $(\mu^C, \mu^P) \in M$ and $\gamma > \bar{\gamma}_2$ for some sufficiently high $\bar{\gamma}_2 < 1$.

Combining the expression for the partial derivatives of $\tilde{\mu}^C$ and $\tilde{\mu}^P$ with these in-

equalities gives

$$\begin{aligned}
\frac{1+\eta}{\lambda}(1-h(\varepsilon, \hat{\mu}^P)) \ln(\beta(\gamma, \varepsilon, h(\varepsilon, \hat{\mu}^P))) &< \frac{\partial \tilde{\mu}^C}{\partial K}(\gamma, \varepsilon, K, L) < 0, \\
\frac{1+\eta}{\lambda}(1-\hat{\mu}^P) \ln(\beta(\gamma, \varepsilon, h(\varepsilon, \hat{\mu}^P))) &< \frac{\partial \tilde{\mu}^P}{\partial K}(\gamma, \varepsilon, K, L) < 0, \\
\frac{1+\eta}{\lambda}(1-h(\varepsilon, \hat{\mu}^P)) \ln(\alpha(\gamma, \varepsilon_C)) &< \frac{\partial \tilde{\mu}^C}{\partial L}(\gamma, \varepsilon, K, L) < 0, \\
\frac{(1+\eta)(\lambda+1)}{\lambda}(1-\hat{\mu}^P) \ln(\alpha(\gamma, \varepsilon_C)) &< \frac{\partial \tilde{\mu}^P}{\partial L}(\gamma, \varepsilon, K, L) < 0,
\end{aligned}$$

for all $\gamma > \max\{\bar{\gamma}_1, \bar{\gamma}_2\}$ and any (K, L) that equals $(\tilde{K}(\gamma, \varepsilon, \mu^C, \mu^P), \tilde{L}(\gamma, \varepsilon, \mu^C, \mu^P))$ for some $(\mu^C, \mu^P) \in M_1 \cap M_2$.

Along with the mean value theorem, these bounds on the partial derivatives of $\tilde{\mu}^C$ and $\tilde{\mu}^P$ imply that there exists $\bar{\gamma} < 1$ and an open neighborhood of $(h(\varepsilon, \hat{\mu}^P), \hat{\mu}^P)$, M , such that

$$\begin{aligned}
0 &\leq \tilde{\mu}^C(\gamma, \varepsilon, \lfloor \tilde{K}(\gamma, \varepsilon, \mu^C, \mu^P) \rfloor, \tilde{L}(\gamma, \varepsilon, \mu^C, \mu^P)) - \mu^C \\
&< -\frac{1+\eta}{\lambda}(1-h(\varepsilon, \hat{\mu}^P)) \ln(\beta(\gamma, \varepsilon, h(\varepsilon, \hat{\mu}^P))), \\
0 &\leq \tilde{\mu}^P(\gamma, \varepsilon, \lfloor \tilde{K}(\gamma, \varepsilon, \mu^C, \mu^P) \rfloor, \tilde{L}(\gamma, \varepsilon, \mu^C, \mu^P)) - \mu^P \\
&< -\frac{1+\eta}{\lambda}(1-\hat{\mu}^P) \ln(\beta(\gamma, \varepsilon, h(\varepsilon, \hat{\mu}^P))),
\end{aligned}$$

for all $\gamma > \bar{\gamma}$ and $(\mu^C, \mu^P) \in M$.

Lemma OA 11 then follows by noting that $\tilde{\mu}^C(\gamma, \varepsilon, \lfloor \tilde{K}(\gamma, \varepsilon, \mu^C, \mu^P) \rfloor, \tilde{L}(\gamma, \varepsilon, \mu^C, \mu^P))$ and $\tilde{\mu}^P(\gamma, \varepsilon, \lfloor \tilde{K}(\gamma, \varepsilon, \mu^C, \mu^P) \rfloor, \tilde{L}(\gamma, \varepsilon, \mu^C, \mu^P))$ is feasible whenever $\tilde{L}(\gamma, \varepsilon, \mu^C, \mu^P)$ is an integer. \blacksquare

Fix $\hat{\mu}^P \in (g/(1+g), 1-g/l)$, $\eta > 0$, and $\lambda > 0$. Let $J_{\hat{\mu}^P, \eta, \lambda}^C : [0, 1] \times (0, 1) \times (0, 1) \times (0, 1) \times (0, 1) \rightarrow \mathbb{R}$ be the function given by

$$\begin{aligned}
&J_{\hat{\mu}^P, \varepsilon, \lambda}^C(\gamma, \varepsilon, \mu^C, \mu^P) \\
&= I\left(1, \varepsilon, \mu^C - \frac{1+\eta}{\lambda}(1-h(\varepsilon, \hat{\mu}^P)) \ln(\beta(\gamma, \varepsilon, h(\varepsilon, \hat{\mu}^P))), \mu^P\right), \tag{OA 10}
\end{aligned}$$

and $J_{\hat{\mu}^P, \eta, \lambda}^D : [0, 1] \times (0, 1) \times (0, 1) \times (0, 1) \times (0, 1) \rightarrow \mathbb{R}$ be the function given by

$$\begin{aligned} & J_{\hat{\mu}^P, \eta, \lambda}^D(\gamma, \varepsilon, \mu^C, \mu^P) \\ &= I\left(\gamma, \varepsilon, \mu^C, \mu^P - \frac{1 + \eta}{\lambda}(1 - \hat{\mu}^P) \ln(\beta(\gamma, \varepsilon, h(\varepsilon, \hat{\mu}^P)))\right). \end{aligned} \quad (\text{OA 11})$$

Lemmas OA 7 and OA 11 imply that, if $|1 + \kappa(\hat{\mu}^P)| > \lambda$, there exists $\bar{\varepsilon} > 0$ such that, for all $\varepsilon_C, \varepsilon_D < \bar{\varepsilon}$ and $\eta > 0$, there exists $\bar{\gamma} < 1$ and an open neighborhood of $(h(\varepsilon, \hat{\mu}^P), \hat{\mu}^P)$, M , such that, for all $\gamma > \bar{\gamma}$ and $(\mu^C, \mu^P) \in M$, whenever $N = \tilde{L}(\gamma, \varepsilon, \mu^C, \mu^P)$ is a non-negative integer, the feasible profile $(\tilde{\mu}^C, \tilde{\mu}^P)$ described in Lemma OA 11 is such that $I(1, \varepsilon, \tilde{\mu}^C, \tilde{\mu}^P) \geq J_{\hat{\mu}^P, \eta, \lambda}^C(\gamma, \varepsilon, \mu^C, \mu^P)$ and $I(\gamma, \varepsilon, \tilde{\mu}^C, \tilde{\mu}^P) \leq J_{\hat{\mu}^P, \eta, \lambda}^D(\gamma, \varepsilon, \mu^C, \mu^P)$.

Next we give conditions under which the γ partial derivatives of $J_{\hat{\mu}^P, \eta, \lambda}^C$ and $J_{\hat{\mu}^P, \eta, \lambda}^D$ evaluated at $(\gamma, \mu^C, \mu^P) = (1, h(\varepsilon, \hat{\mu}^P), \hat{\mu}^P)$ are both strictly negative, and are such that the γ partial derivative of $J_{\hat{\mu}^P, \eta, \lambda}^D$ is strictly less than that of $J_{\hat{\mu}^P, \eta, \lambda}^C$. An implication of this is that, for all sufficiently high γ , there is a (μ^C, μ^P) isocurve of $I(1, \varepsilon, \mu^C, \mu^P)$ in M such that $J_{\hat{\mu}^P, \eta, \lambda}^D(\gamma, \varepsilon, \mu^C, \mu^P) < 0 < J_{\hat{\mu}^P, \eta, \lambda}^C(\gamma, \varepsilon, \mu^C, \mu^P)$ for all (μ^C, μ^P) on the isocurve.

Lemma OA 12. *Fix $\hat{\mu}^P \in (g/(1+g), 1-g/l]$. If there is λ satisfying $|1 + \kappa(\hat{\mu}^P)| > \lambda > \iota(\hat{\mu}^P)$, then there is $\eta > 0$ and $\bar{\varepsilon} > 0$ such that, for all $\varepsilon_C, \varepsilon_D < \bar{\varepsilon}$,*

$$0 < \frac{\partial J_{\hat{\mu}^P, \eta, \lambda}^C}{\partial \gamma}(1, \varepsilon, h(\varepsilon, \hat{\mu}^P), \hat{\mu}^P) < \frac{\partial J_{\hat{\mu}^P, \eta, \lambda}^D}{\partial \gamma}(1, \varepsilon, h(\varepsilon, \hat{\mu}^P), \hat{\mu}^P).$$

Proof of Lemma OA 12. Differentiating Equation OA 10, we find that

$$\begin{aligned} & \frac{\partial J_{\hat{\mu}^P, \eta, \lambda}^C}{\partial \gamma}(1, \varepsilon, h(\varepsilon, \hat{\mu}^P), \hat{\mu}^P) \\ &= -\frac{1 + \eta}{\lambda} \frac{1 - h(\varepsilon, \hat{\mu}^P)}{1 - \varepsilon_D - (1 - \varepsilon_C - \varepsilon_D)h(\varepsilon, \hat{\mu}^P)} \frac{\partial I}{\partial \mu^C}(1, \varepsilon, h(\varepsilon, \hat{\mu}^P), \hat{\mu}^P) \\ &= \frac{l(1 - \hat{\mu}^P)}{1 - \varepsilon_D - (1 - \varepsilon_C - \varepsilon_D)h(\varepsilon, \hat{\mu}^P)} \frac{1 + \eta}{\lambda} \left(\frac{\varepsilon_C}{1 - \varepsilon_D - (1 - \varepsilon_C - \varepsilon_D)h(\varepsilon, \hat{\mu}^P)} \right. \\ & \quad \left. - \frac{1 - h(\varepsilon, \hat{\mu}^P)}{l(1 - \hat{\mu}^P)} \rho(1, \varepsilon, h(\varepsilon, \hat{\mu}^P))(l - g) \right). \end{aligned}$$

Differentiating Equation OA 11, we find that

$$\begin{aligned}
& \frac{\partial J_{\hat{\mu}^P, \eta, \lambda}^D}{\partial \gamma}(1, \varepsilon, h(\varepsilon, \hat{\mu}^P), \hat{\mu}^P) \\
&= \frac{\partial I}{\partial \gamma}(1, \varepsilon, h(\varepsilon, \hat{\mu}^P), \hat{\mu}^P) \\
&\quad - \frac{1 + \eta}{\lambda} \frac{1 - \hat{\mu}^P}{1 - \varepsilon_D - (1 - \varepsilon_C - \varepsilon_D)h(\varepsilon, \hat{\mu}^P)} \frac{\partial I}{\partial \mu^P}(1, \varepsilon, h(\varepsilon, \hat{\mu}^P), \hat{\mu}^P) \\
&= \frac{l(1 - \hat{\mu}^P)}{1 - \varepsilon_D - (1 - \varepsilon_C - \varepsilon_D)h(\varepsilon, \hat{\mu}^P)} \left(1 - \frac{1 + \eta}{\lambda} \left(\frac{1}{l} \rho(1, \varepsilon, h(\varepsilon, \hat{\mu}^P))(1 + g - l) + 1 \right) \right).
\end{aligned}$$

Note that

$$\begin{aligned}
& \lim_{\varepsilon \rightarrow 0} \frac{1 - \varepsilon_D - (1 - \varepsilon_C - \varepsilon_D)h(\varepsilon, \hat{\mu}^P)}{l(1 - \hat{\mu}^P)} \frac{\partial J_{\hat{\mu}^P, \eta, \lambda}^C}{\partial \gamma}(1, \varepsilon, h(\varepsilon, \hat{\mu}^P), \hat{\mu}^P) \\
&= \frac{1 + \eta}{\lambda} \frac{(1 + g)\hat{\mu}^P - g}{(1 + g - l)\hat{\mu}^P + l - g}
\end{aligned}$$

and

$$\begin{aligned}
& \lim_{\varepsilon \rightarrow 0} \frac{1 - \varepsilon_D - (1 - \varepsilon_C - \varepsilon_D)h(\varepsilon, \hat{\mu}^P)}{l(1 - \hat{\mu}^P)} \frac{\partial J_{\hat{\mu}^P, \eta, \lambda}^D}{\partial \gamma}(1, \varepsilon, h(\varepsilon, \hat{\mu}^P), \hat{\mu}^P) \\
&= 1 - \frac{1 + \eta}{\lambda} \left(\frac{1}{(1 + g - l)\hat{\mu}^P + l - g} \right).
\end{aligned}$$

When $\lambda > \iota(\hat{\mu}^P)$,

$$1 - \frac{1}{\lambda} \left(\frac{1}{(1 + g - l)\hat{\mu}^P + l - g} \right) > \frac{1}{\lambda} \frac{(1 + g)\hat{\mu}^P - g}{(1 + g - l)\hat{\mu}^P + l - g} > 0,$$

so there is $\eta > 0$ such that

$$1 - \frac{1 + \eta}{\lambda} \left(\frac{1}{(1 + g - l)\hat{\mu}^P + l - g} \right) > \frac{1 + \eta}{\lambda} \frac{(1 + g)\hat{\mu}^P - g}{(1 + g - l)\hat{\mu}^P + l - g} > 0.$$

Thus, for such an η , there exists $\bar{\varepsilon}$ such that

$$0 < \frac{\partial J_{\hat{\mu}^P, \eta, \lambda}^C}{\partial \gamma}(1, \varepsilon, h(\varepsilon, \hat{\mu}^P), \hat{\mu}^P) < \frac{\partial J_{\hat{\mu}^P, \eta, \lambda}^D}{\partial \gamma}(1, \varepsilon, h(\varepsilon, \hat{\mu}^P), \hat{\mu}^P)$$

for all $\varepsilon_C, \varepsilon_D < \bar{\varepsilon}$. ■

Lemma OA 13. Fix $\hat{\mu}^P \in (g/(1+g), 1-g/l]$. There is $\bar{\varepsilon} > 0$ such that, for all $\varepsilon_C, \varepsilon_D < \bar{\varepsilon}$, the isocurves of $\tilde{L}(\gamma, \varepsilon, \mu^C, \mu^P)$ and $I(1, \varepsilon, \mu^C, \mu^P)$ are not tangent at $(h(\varepsilon, \hat{\mu}^P), \hat{\mu}^P)$.

Proof of Lemma OA 13. By Lemma OA 8, the isocurve of $\tilde{L}(\gamma, \varepsilon, \mu^C, \mu^P)$ has slope

$$\begin{aligned} \frac{d\mu^C}{d\mu^P} &= -\frac{\frac{\partial \tilde{L}}{\partial \mu^P}(\gamma, \varepsilon, h(\varepsilon, \hat{\mu}^P), \hat{\mu}^P)}{\frac{\partial \tilde{L}}{\partial \mu^C}(\gamma, \varepsilon, h(\varepsilon, \hat{\mu}^P), \hat{\mu}^P)} \\ &= \frac{1 - h(\varepsilon, \hat{\mu}^P)}{1 - \hat{\mu}^P} \end{aligned}$$

at $(h(\varepsilon, \hat{\mu}^P), \hat{\mu}^P)$.

Likewise, we find that the isocurve of $I(1, \varepsilon, \mu^C, \mu^P)$ has slope

$$\begin{aligned} \frac{d\mu^C}{d\mu^P} &= -\frac{\frac{\partial I}{\partial \mu^P}(1, \varepsilon, h(\varepsilon, \hat{\mu}^P), \hat{\mu}^P)}{\frac{\partial I}{\partial \mu^C}(1, \varepsilon, h(\varepsilon, \hat{\mu}^P), \hat{\mu}^P)} \\ &= \left(\frac{\rho(1, \varepsilon, h(\varepsilon, \hat{\mu}^P))(1+g-l) + l}{\frac{\varepsilon_C}{1-\varepsilon_D-(1-\varepsilon_C-\varepsilon_D)h(\varepsilon, \hat{\mu}^P)}l - \frac{1-h(\varepsilon, \hat{\mu}^P)}{1-\hat{\mu}^P}\rho(1, \varepsilon, h(\varepsilon, \hat{\mu}^P))(l-g)} \right) \\ &\quad \left(\frac{1-h(\varepsilon, \hat{\mu}^P)}{1-\hat{\mu}^P} \right) \end{aligned}$$

at $(h(\varepsilon, \hat{\mu}^P), \hat{\mu}^P)$.

Since

$$\begin{aligned} &\lim_{\varepsilon \rightarrow 0} \frac{\rho(1, \varepsilon, h(\varepsilon, \hat{\mu}^P))(1+g-l) + l}{\frac{\varepsilon_C}{1-\varepsilon_D-(1-\varepsilon_C-\varepsilon_D)h(\varepsilon, \hat{\mu}^P)}l - \frac{1-h(\varepsilon, \hat{\mu}^P)}{1-\hat{\mu}^P}\rho(1, \varepsilon, h(\varepsilon, \hat{\mu}^P))(l-g)} \\ &= \frac{1}{(1+g)\hat{\mu}^P - g} > 1, \end{aligned}$$

the result follows. ■

Combining Lemmas OA 12 and OA 13 gives the following result.

Lemma OA 14. Fix $\hat{\mu}^P \in (g/(1+g), 1-g/l]$. If $|1 + \kappa(\hat{\mu}^P)| > \iota(\hat{\mu}^P)$, there exists $\bar{\varepsilon} > 0$ such that, for all $\varepsilon_C, \varepsilon_D < \bar{\varepsilon}$ and all open neighborhoods of $(h(\varepsilon, \hat{\mu}^P), \hat{\mu}^P)$, M ,

there exists $\bar{\gamma} < 1$ such that, for all $\gamma > \bar{\gamma}$, there is a feasible $(\mu^C, \mu^P) \in M$ that satisfies the incentive constraints.

Proof of Lemma OA 14. By Lemma OA 12, there exists $\bar{\gamma} < 1$, sufficiently small neighborhood of $(\mu^C, \mu^P) = (h(\varepsilon, \hat{\mu}^P), \hat{\mu}^P)$, M , and $\eta_1, \eta_2 > 0$ such that

$$0 < \frac{\partial J_{\hat{\mu}^P, \eta, \lambda}^C}{\partial \gamma}(\gamma, \varepsilon, \mu^C, \mu^P) < \eta_1 < \eta_2 < \frac{\partial J_{\hat{\mu}^P, \eta, \lambda}^D}{\partial \gamma}(\gamma, \varepsilon, \mu^C, \mu^P)$$

for all $(\mu^C, \mu^P) \in M$ and $\gamma > \bar{\gamma}$. Therefore,

$$\begin{aligned} J_{\hat{\mu}^P, \eta, \lambda}^C(\gamma, \varepsilon, \mu^C, \mu^P) &\geq J_{\hat{\mu}^P, \eta, \lambda}^C(1, \varepsilon, \mu^C, \mu^P) - \eta_1(1 - \gamma) \\ &= I(1, \varepsilon, \mu^C, \mu^P) - \eta_1(1 - \gamma) \\ J_{\hat{\mu}^P, \eta, \lambda}^D(\gamma, \varepsilon, \mu^C, \mu^P) &\leq J_{\hat{\mu}^P, \eta, \lambda}^D(1, \varepsilon, \mu^C, \mu^P) - \eta_2(1 - \gamma) \\ &= I(1, \varepsilon, \mu^C, \mu^P) - \eta_2(1 - \gamma) \end{aligned}$$

for all $(\mu^C, \mu^P) \in M$ and $\gamma > \bar{\gamma}$. It thus follows that if there is some $(\mu^C, \mu^P) \in M$ such that $\tilde{L}(\gamma, \varepsilon, \mu^C, \mu^P)$ is a non-negative integer and that satisfies $\eta_1(1 - \gamma) < I(1, \varepsilon, \mu^C, \mu^P) < \eta_2(1 - \gamma)$ and $\mu^P \leq 1 - g/l$, then $(\tilde{\mu}^C(\gamma, \varepsilon, \mu^C, \mu^P), \tilde{\mu}^P(\gamma, \varepsilon, \mu^C, \mu^P))$ is both feasible and satisfies all of the incentive constraints for γ .

All that remains is to show that, for all $\gamma > \bar{\gamma}$, there exists $(\mu^C, \mu^P) \in M$ for which these conditions are met. Because

$$\frac{\partial I}{\partial \mu^C}(1, \varepsilon, h(\varepsilon, \hat{\mu}^P), \hat{\mu}^P) < 0,$$

it follows that, for sufficiently large γ , isocurves of the form $I(1, \varepsilon, \mu^C, \mu^P) = (\eta_1 + \eta_2)/2(1 - \gamma)$ intersect M for every μ^P in an open neighborhood of $1 - g/l$. By Lemma OA 13, the isocurves of $I(1, \varepsilon, \mu^C, \mu^P)$ and $\tilde{L}(\gamma, \varepsilon, \mu^C, \mu^P)$ are not tangent. Because the $\tilde{L}(\gamma, \varepsilon, \mu^C, \mu^P)$ isocurves do not depend on γ and

$$\lim_{\gamma \rightarrow 1} \ln(\alpha(\gamma, \varepsilon_C)) = 0,$$

it follows by Lemma OA 8 that there exists $(\mu^C, \mu^P) \in M$ on the isocurve $I(1, \varepsilon, \mu^C, \mu^P) = (\eta_1 + \eta_2)/2(1 - \gamma)$ that satisfies $\mu^P \leq 1 - g/l$ and is such that $\tilde{L}(\gamma, \varepsilon, \mu^C, \mu^P)$ is a non-negative integer for sufficiently large γ . \blacksquare

Lemma OA 5 is an immediate consequence of Lemma OA 14.

OA.4.2 Proof of Lemma OA 6

Lemma OA 6 is a consequence of the following claim.

Claim OA 1. *Some $\mu^P \in (g/(1+g), 1 - g/l]$ satisfies $|1 + \kappa(\mu^P)| > \iota(\mu^P)$ if any of the following conditions hold.*

1. $g < e - 1$ and

$$l > \frac{1+g}{1 - \ln(1+g)}.$$

2. $g > e - 1$ and

$$l > \frac{1+g}{\ln(1+g) - 1}.$$

3. For some $\phi > 1$, $g < e^\phi - 1$, $l \geq e^\phi g$, and

$$l > \frac{3e^\phi - 2 - 2g}{\phi - 1}.$$

To see why this implies Lemma OA 6, note that if $l > (1+g)/|\ln(1+g) - 1|$ then either the first or second condition of Claim OA 1 is satisfied. Moreover, if $l > g + g^2$ and $l > 22.0625 - 3.63636g$ (which is just $l > \frac{3e^\phi - 2 - 2g}{\phi - 1}$ for $\phi = 1.55$), but $l \leq (1+g)/|\ln(1+g) - 1|$, then it must be that $g < 2.641 < e^{1.55} - 1$. Since $22.0625 - 3.63636g > e^{1.55}g$ for all such g , it follows that the third condition of Claim OA 1 is satisfied in this case.

Proof of Claim OA 1. We handle Cases 2 and 3. The proof for Case 1 is similar to that for Case 2.

Suppose that $g > e - 1$ and $l > (1 + g)/(\ln(1 + g) - 1)$. Note that

$$\lim_{\mu^P \rightarrow \frac{g}{1+g}} |1 + \kappa(\mu^P)| - \iota(\mu^P) = \ln(1 + g) - 1 - \frac{1 + g}{l}.$$

Since $l > (1 + g)/(\ln(1 + g) - 1)$, $\ln(1 + g) - 1 - (1 + g)/l > 0$, and the result follows.

Suppose that, for some $\phi > 1$, $g < e^\phi - 1$, $l \geq e^\phi g$ and $l > (3e^\phi g - 2 - 2g)/(\phi - 1)$. Note that $g/(1 + g) < 1 - e^{-\phi} \leq 1 - g/l$ and

$$|1 + \kappa(1 - e^{-\phi})| - \iota(1 - e^{-\phi}) = \frac{|l(\phi - 1) - e^\phi + 1 + g| - 2e^\phi + 1 + g}{e^\phi - 1 - g + l}.$$

Since $l > (3e^\phi - 2 - 2g)/(\phi - 1)$, $|l(\phi - 1) - e^\phi + 1 + g| - 2e^\phi + 1 + g > 0$, and the result follows. ■

OA.5 Limit-Supporting C with PPR

OA.5.1 Limit-Supporting C for $g < 1$

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

Lemma OA 15. *There is a $P_{\chi_P} D$ equilibrium with share μ^C of players in P iff the following feasibility constraint*

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

and incentive constraint

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

are satisfied.

Proof of Lemma OA 15. The steady-state condition for μ^P is $1 - \gamma = (1 - \gamma)\mu^P + \gamma\chi_P(1 - \varepsilon_D - (1 - \varepsilon_C - \varepsilon_D)\mu^C)\mu^P$, which is equivalent to $\chi_P = \frac{(1-\gamma)(1-\mu^P)}{\gamma(1-\varepsilon_D-(1-\varepsilon_C-\varepsilon_D)\mu^C)\mu^P}$.

Since V_0 equals average payoff in the population in every period, the value function for a preciprocator is $V^P = (\mu^C)^2$. Thus,

$$\frac{1 - \varepsilon_C - \varepsilon_D}{1 - \varepsilon_D - (1 - \varepsilon_C - \varepsilon_D)\mu^C}(\mu^C - (\mu^C)^2) = \frac{(1 - \varepsilon_C - \varepsilon_D)(1 - \mu^C)}{1 - \varepsilon_D - (1 - \varepsilon_C - \varepsilon_D)\mu^C}\mu^C.$$

Combining this with Lemma 8 and the fact that $p_k^D = 1 - \varepsilon_D - (1 - \varepsilon_C - \varepsilon_D)\mu^C$ for any preciprocator record, the incentives of the preciprocators are satisfied iff $g < \frac{(1-\varepsilon_C-\varepsilon_D)(1-\mu^C)}{1-\varepsilon_D-(1-\varepsilon_C-\varepsilon_D)\mu^C}\mu^C < l$. ■

We now show that C is supportable in the general $(\gamma, \varepsilon) \rightarrow (1, 0)$ limit whenever $g < 1$.

Lemma OA 16. *If $g < 1$, then there is a sequence of equilibria in which $\lim_{(\gamma, \varepsilon) \rightarrow (1, 0)} \mu^C(\gamma, \varepsilon) = 1$.*

Proof of Lemma OA 16. Fix some κ such that $g < \kappa/(1 + \kappa) < \min\{l, 1\}$. Consider supporting $\mu^C(\gamma, \varepsilon) = \tilde{\mu}^C(\varepsilon_C) = 1 - \kappa\varepsilon_C$. Note that $\tilde{\mu}^C(\varepsilon)$ is well-defined for sufficiently small ε and satisfies

$$\lim_{\varepsilon \rightarrow 0} \frac{(1 - \varepsilon_C - \varepsilon_D)(1 - \tilde{\mu}^C(\varepsilon_C))}{1 - \varepsilon_D - (1 - \varepsilon_C - \varepsilon_D)\tilde{\mu}^C(\varepsilon_C)}\tilde{\mu}^C(\varepsilon_C) = \frac{\kappa}{1 + \kappa}.$$

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

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

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

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

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

OA.5.2 Limit-Supporting C for $l > g + g^2$

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

Lemma OA 17. *There is a $D_{\chi_{D_1}} P_{\chi_P} S_{\chi_S} D$ equilibrium with share μ^{D_1} of players in D_1 , share μ^P of players in P , share μ^S of players in S , and share μ^{D_2} of players in D_2 iff the following feasibility constraints*

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

and incentive constraints

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

are satisfied.

Proof of Lemma OA 17. The steady-state condition for μ^{D_1} is equivalent to $\chi_{D_1} =$

$\frac{(1-\gamma)(1-\mu^{D_1})}{\gamma(1-\varepsilon_D)\mu^{D_1}}$. Likewise, the steady-state condition for μ^P is

$$\gamma(1-\varepsilon_D)\chi_{D_1}\mu^{D_1} = (1-\gamma)\mu^P + \gamma\chi_P(1-\varepsilon_D - (1-\varepsilon_C - \varepsilon_D)\mu^C)\mu^P,$$

which is equivalent to

$$\chi_P = \frac{(1-\gamma)(1-\mu^{D_1} - \mu^P)}{\gamma(1-\varepsilon_D - (1-\varepsilon_C - \varepsilon_D)\mu^C)\mu^P}.$$

Finally, the steady-state condition for μ^S is

$$\gamma\chi_P(1-\varepsilon_D - (1-\varepsilon_C - \varepsilon_D)\mu^C)\mu^P = (1-\gamma)\mu^S + \gamma\chi_S\varepsilon_C\mu^S,$$

which is equivalent to

$$\chi_S = \frac{(1-\gamma)(1-\mu^{D_1} - \mu^P - \mu^S)}{\gamma\varepsilon_C\mu^S} = \frac{(1-\gamma)\mu^{D_2}}{\gamma\varepsilon_C\mu^S}.$$

The value function for a reciprocator is

$$V^P = \frac{1}{1-\mu^{D_1}}(\mu^P\mu^C + \mu^S(\mu^C - \mu^{D_1}) + \mu^D\mu^S(1+g)),$$

since V^P equals the average payoff of players in the population who are in either P , S , or D_2 . in every period. Thus,

$$\begin{aligned} \frac{1-\varepsilon_C - \varepsilon_D}{1-\varepsilon_D - (1-\varepsilon_C - \varepsilon_D)\mu^C}(\mu^C - V^P) &= \frac{1-\varepsilon_C - \varepsilon_D}{1-\varepsilon_D - (1-\varepsilon_C - \varepsilon_D)\mu^C} \\ &\quad \left(\mu^C - \frac{1}{1-\mu^{D_1}}(\mu^P\mu^C + \mu^S(\mu^C - \mu^{D_1}) + \mu^D\mu^S(1+g)) \right) \\ &= \frac{1}{1-\mu^{D_1}} \frac{(1-\varepsilon_C - \varepsilon_D)(1-\mu^C)}{1-\varepsilon_D - (1-\varepsilon_C - \varepsilon_D)\mu^C} \\ &\quad \left(\mu^C + \mu^S l - \frac{\mu^{D_1}}{\mu^D}\mu^C - \frac{\mu^{D_2}}{\mu^D}\mu^S(1+g) \right). \end{aligned}$$

Combining this with Lemma 8 and the fact that $p_k^D = 1 - \varepsilon_D - (1 - \varepsilon_C - \varepsilon_D)\mu^C$ for

any preciprocator record, the incentives of the preciprocators are satisfied iff

$$g < \frac{1}{1 - \mu^{D_1}} \frac{(1 - \varepsilon_C - \varepsilon_D)(1 - \mu^C)}{1 - \varepsilon_D - (1 - \varepsilon_C - \varepsilon_D)\mu^C} \left(\mu^C + \mu^S l - \frac{\mu^{D_1}}{\mu^D} \mu^C - \frac{\mu^{D_2}}{\mu^D} \mu^S (1 + g) \right) < l.$$

Moreover, the value function for a supercooperator is

$$V^S = \frac{1}{1 - \mu^{D_1} - \mu^P} (\mu^S (\mu^C - \mu^D l) + \mu^D \mu^S (1 + g)),$$

since V^S equals the average payoff of players in the population who are in either S or D_2 . Thus,

$$\begin{aligned} \frac{1 - \varepsilon_C - \varepsilon_D}{\varepsilon_C} (\mu^C - \mu^D l - V^S) &= \frac{1 - \varepsilon_C - \varepsilon_D}{\varepsilon_C} \\ &\quad \left(\mu^C - \mu^D l - \frac{1}{1 - \mu^{D_1} - \mu^P} (\mu^S (\mu^C - \mu^D l) + \mu^D \mu^S (1 + g)) \right) \\ &= \frac{1 - \varepsilon_C - \varepsilon_D}{\varepsilon_C} \frac{\mu^{D_2}}{\mu^S + \mu^{D_2}} (\mu^C - \mu^D l - \mu^S (1 + g)). \end{aligned}$$

Combining this with Lemma 8 and the fact that $p_k^D = \varepsilon_C$ for any supercooperator record, the incentives of the supercooperators are satisfied iff $\frac{1 - \varepsilon_C - \varepsilon_D}{\varepsilon_C} \frac{\mu^{D_2}}{\mu^S + \mu^{D_2}} (\mu^C - \mu^D l - \mu^S (1 + g)) > l$. \blacksquare

We now show that we can support C in the general $(\gamma, \varepsilon) \rightarrow (1, 0)$ limit whenever $l > g + g^2$. Consider the following inequalities:

$$\begin{aligned} \mu^C + \mu^S (l - g - 1) &> g, \\ \mu^C - \mu^D l - \mu^S (1 + g) &> 0. \end{aligned} \tag{OA 12}$$

Let $A(g, l)$ be the set of $\mu^C \in (0, 1)$ such that there exists some $g/l < \mu^S < \mu^C$ where the above inequalities are satisfied. We first establish that, for any $\mu^C \in A(g, l)$, there is a sequence of equilibria in which the share of cooperators converges to μ^C as $(\gamma, \varepsilon) \rightarrow (1, 0)$. We then show that $A(g, l)$ contains values of μ^C that are arbitrarily close to 1, which completes the proof that C can be limit-supported.

Lemma OA 18. Fix some $\mu^C \in A(g, l)$. If $l > g + g^2$, there is a sequence of equilibria in which $\lim_{(\gamma, \varepsilon) \rightarrow (1, 0)} \mu^C(\gamma, \varepsilon) = \mu^C$.

Proof of Lemma OA 18. Fix some $g/l < \mu^S < \mu^C$ such that the inequalities in (OA 12) are satisfied. Consider supporting such μ^C and μ^S along with

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

where $0 < \lambda < 1$ is taken to be sufficiently small so that $\mu^P + (1 + \lambda)\mu^S < 1$. We will argue that for all (γ, ε) sufficiently close to $(1, 0)$, such shares can be supported in equilibrium.

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

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

Additionally,

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

and

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

for all (γ, ε) .

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

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

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

$(1 - \varepsilon_C - \varepsilon_D)\mu^C = 1$, it follows that

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

for sufficiently small ε_C and ε_D . Moreover,

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

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

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

for all (γ, ε) .

Furthermore,

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

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

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

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

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

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

for sufficiently small $\kappa_2 > 0$,

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

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