Robust Cooperation with First-Order Information*

Daniel Clark, Drew Fudenberg, and Alexander Wolitzky

April 19, 2019

Abstract

We study the repeated prisoner's dilemma in a large population with random matching and overlapping generations. Our goal is to study the extent to which cooperation can be supported by equilibria where players use simple strategies and have very limited information about their opponents. To this end, we assume players have access to only *first-order information* about their current partners, meaning that a player's record tracks information about her past actions only, and not her partners' past actions (or her partners' partners' actions). Cooperation in strict equilibrium is impossible if payoffs are submodular, or if players can erase their records. If payoffs are sufficiently supermodular, then cooperation can be sustained by a tolerant version of grim trigger strategies, where a player retains good standing until her record reflects a certain number of defections. Players close to the threshold may cooperate even with defectors; this "unconditional cooperation" is crucial for sustaining maximal cooperation. If players can forge records of fake past interactions and successfully coordinate with their current partners, this tolerant version of grim trigger is the unique family of strategies that sustains cooperation in equilibrium.

*We thank V Bhaskar, Glenn Ellison, 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 finanical support.

1 Introduction

In many settings of economic interest, individuals interact with different partners over time, and bad behaviour against one partner causes a negative response by other members of society. This paper studies such "community enforcement" in the repeated prisoner's dilemma. Our two goals are to better understand what sorts of information suffice for good community outcomes and to provide a foundation for the analysis of relatively simple strategies that we think might be descriptively plausible. To this end, we depart from past models of community enforcement (which we review in Section 2) by dropping the assumption of a common start date and calendar time, and by considering quite restrictive forms of information and record-keeping. We then study when cooperation is possible under these restrictions, and how the maximal level of cooperation varies with the parameters of the underlying game.

To place our work in context, recall that in the standard repeated game model, a fixed set of players interacts repeatedly with a commonly known start date and a common notion of calendar time. When each player's signals are sufficient to statistically identify the vector of her 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 number of relatively sophisticated players, such as business partnerships or collusive agreements among firms. However, laboratory studies of repeated games suggest that many subjects use fairly simple strategies.[1] Moreover, repeated games have also been used to model cooperation in large populations, and for these applications the assumptions of a fixed population, a common start date, and common calendar time seem less appropriate, and there is all the more reason to doubt whether players will use complicated strategies.

Thus, instead of analyzing interactions in fixed pairs, we consider a class of repeated

[1]See e.g. Fudenberg, Rand, and Dreber (2012) and the survey by Dal Bó and Fréchette (2018).

games with random matching, no commonly known start date or calendar time, and restrictive forms of information about past play. In our model, there is a continuum of players with geometrically distributed lifespans. Each player has a series of pairwise interactions with different partners. We assume players have only "first-order" information about their partners, meaning that their information depends only on the actions the partner has taken, and not on the actions or histories of the partner's past partners. In particular, each player carries a "record" which depends only on her own past actions (perhaps stochastically), and when two players meet they observe each other's record and nothing else. We study steady states of this population game, where the share of the population with each record is constant over time, and each player's strategy depends only on her own record and the record of her current opponent. A preliminary result is that a steady state always exists.

We then assume the stage game is the standard prisoner's dilemma:

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

Figure 1: The Prisoner's Dilemma

with $g, l > 0$ and $l + 1 > g$, so (C,C) maximizes the sum of payoffs. Here g measures the gain to defection (that is, playing D) when one's opponent cooperates, for example the cost savings from providing a low quality product or service, or the profit gained by undercutting competitors in a cartel. Moreover, because l measures the gain from playing D against D, the comparison of g and l is a measure of the complementarity in the interaction. As we will see, the possibility and maximal extent of equilibrium cooperation in our setting depend on the degree of complementarity as well as the temptation to deviate. Intuitively, this is because first-order information is not sufficient to distinguish between opportunistic deviations to D and equilibrium plays of D that punish opponents with bad records, and so players must sometimes be willing to worsen their record by playing D against D when their continuation payoff would be higher if they played C and incurred a short run loss.

Throughout the paper we restrict attention to strict equilibria; this captures a simple form of robustness and, in particular, rules out "belief-free" equilibria and related constructions. The steady state where everyone always plays D regardless of the records is always a strict equilibrium. As in the related random matching models of Takahashi (2010) and Heller and Mohlin (2018), we find that when the prisoner's dilemma stage game is "submodular," that is when $g \geq l$, the only strict equilibrium is *Always Defect*, regardless of the community's record-keeping system. We thus focus on the supermodular case where $g < l$. In our setting, though, this is not sufficient to permit cooperative equilibria; we will see that we additionally require sufficiently high complementarity.

The space of possible record-keeping systems that one could consider is very large. To sharpen our focus, we note that in some economic environments players may be able to manipulate their records, most notably by evading record-keeping or by forging fake positive records. We show that, if players can evade record-keeping, the only strict equilibrium is again *Always Defect*. We thus consider record-keeping systems where evasion is impossible, and instead ask for robustness to forging records of fake interactions, a property we call *forgery-proofness*.[2] This consideration, along with an interest in simplicity, leads us to study the very simple technology that records only the number of times a player has played D, but not when those times were or how many times the player has played C (as rewarding players for having more C's in their record would encourage them to fake C's).

For this information structure, we study how much cooperation can be supported in

[2]Fake reviews are an important problem for many real-world reputation systems. For example, fake reviews on Yelp are sufficiently common that Yelp uses an algorithm to (imperfectly) filter many of them out (Luca and Zervas, 2016). For an empirical analysis of review manipulation on Expedia and TripAdvisor, see Mayzlin, Dover, and Chevalier (2014). In the context of online book reviews, Chevalier and Mayzlin (2006) find that the relatively rare one-star reviews carry a lot of weight with consumers. They argue that this is because "the author can post a large number of meaningless five-star reviews cheaply, (but) cannot prevent others from posting one-star reviews," (pp. 349-350).

the limit as expected lifespans grow to infinity (i.e. the continuation probability $\gamma \to 1$) and "noise" (i.e., the ε probability that a play of C is mis-recorded as D) shrinks to 0, and say that there is a *limit efficient* equilibrium when this share converges to 1. Our first observation here is that limit efficient strict equilibria always exist if the prisoner's dilemma is *mild*, meaning that $g < 1$. The strategies we use in this construction are "forgiving," in the sense that no matter how many D's there are in a player's record, the appropriate sequence of play will lead to them being forgiven and treated as a cooperator. Moreover, as in the strategy "perfect tit-for-tat" (Fudenberg and Maskin, 1990), the way a player earns forgiveness is by acquiring additional D's in her record. Thus, a player with more D's in her record is sometimes treated better than a player with fewer D's. For this reason, the equilibrium is not forgery-proof.

However, forgery-proofness on its own does not rule out all the equilibria that we find implausible. In particular, it allows there to be multiple disjoint sets of records, where players with records in each set cooperate with each other but not with the reciprocators in other sets: For example, players with 2 D's might cooperate only with other players with 2 D's, and players with 4 D's might only cooperate with other players whose record is 4. Thus we also require that strategies are *coordination-proof*, which means that, if the prisoner's dilemma faced by a pair of players becomes a coordination game when it is augmented by their equilibrium continuation payoffs (as a function of their records), they play the Pareto-dominant equilibrium. Combined, these restriction imply that each player uses a strategy of a form we call $GrimKL$. This strategy partitions players into three groups: players with $K-1$ or fewer D's in their records cooperate with any player with $K+L-1$ or fewer D's; players with between K and $K+L-1$ D's are unconditional cooperators (who cooperate with everyone); and players with $K+L$ or more D's are unconditional defectors. We call the special case where $L=0$—so there are no unconditional cooperators—$GrimK$.[3]

[3] $GrimK$ was introduced by Fudenberg, Rand, and Dreber (2012) who noted that while, for $K > 1$, it is never an equilibrium in a two player game with perfect monitoring, it can be an equilibrium when actions are observed with noise, and that some experimental subjects seem to use such strategies

Having identified *GrimK* and *GrimKL* as the only plausible strategies in our environment, we then study when these strategies can support cooperation, and in particular ask when they are limit efficient (for appropriately chosen K and L). We find that *GrimK* cannot support any cooperation when the static incentive to deviate from cooperation is "large" in the sense that $g \geq 1$, regardless of the value of the continuation probability γ. When instead $g < 1$ and $l > g/(1-g)$, *GrimK* can support some cooperation in the limit where $\gamma \to 1$ and then $\varepsilon \to 0$, and the degree of limit inefficiency shrinks as l grows. To see the intuition, note that, since a player gains at least g by deviating from C to D, a player's continuation payoff must be reduced by g whenever her record acquires a D. With *GrimK* strategies, this reduction must be achieved by "switching" g future plays of the stage game from (C, C) to (D, D). Since each such switch also adds a D to the partner's record, this leads to additional switches. When $g < 1$, these switches dampen over time, while when $g \geq 1$ they "snowball." Thus, positive steady-state cooperation is possible when $g < 1$, but not when $g \geq 1$.[4]

Our last main result is that the more flexible *GrimKL* strategies cannot support cooperation in the limit when $l < g(1+g)$ (regardless of γ) but are limit efficient when $l > g(1+g)$.[5] (The condition $g \gtrless 1$ plays no role here.) Thus, unlike *GrimK* strategies, *GrimKL* strategies are sometimes fully limit efficient, and they can be limit efficient even for parameters where *GrimK* strategies cannot support any cooperation. Specifically, *GrimKL* does not require a "small" static gain from defection, and requires less complementarity than *GrimK* does. The reason *GrimKL* strategies can be effective even when $g \geq 1$ is that they incorporate "extra effort" in addition to "punishment": With *GrimK* strategies a player's continuation payoff can fall only by switching play from (C, C) to (D, D), while with *GrimKL* strategies a player's continuation payoff can also fall by switching play from (C, C) to (C, D). This additional loss makes it possible to avoid the "snowballing" of D's that precludes cooperation

[4]Moreover, as we show in section 5.1, this result extends to any strategy that does not involve unconditional cooperators.

[5]Technically, the lower bound on l required for our limit efficiency result may be slightly greater than $g(1+g)$ due to an integer problem. See Section 6.3.

with $GrimK$ strategies when $g \geq 1$.

We find it particularly interesting that $GrimKL$ strategies can support cooperation even in severe prisoner's dilemmas ($g \geq 1$) by using extra effort from players with marginal records to cut off the snowballing defections that would arise under $GrimK$. This is somewhat related to notions of "repentance" or "restitution" that have been shown to support sustainable cooperation in some well-known case studies (Milgrom, North, and Weingast, 1990; Ellickson, 1994). However, under repentance a player with marginal standing can improve her standing by cooperating, while under $GrimKL$ the best she can do is prevent further deterioration of her status. This difference comes from the fact that we want equilibria to be forgery-proof: if players can forge records that say they cooperated, repentance strategies break down, while $GrimKL$ is robust.

All our main results involve records that (noisily) track the number of times a player has played D. If records instead track the number of times a player has played C, we show cooperation is impossible. More powerful record-keeping systems—for example, systems that track the number of both C's and D's, or that track the timing of actions—could sometimes support more cooperation. However, exploiting these systems may require more complicated strategies, and we show that (very simple) $GrimKL$ strategies can already attain full limit efficiency when only D's are tracked, as long as the game is "sufficiently supermodular." Nonetheless, investigating different information structures is one promising direction for future research.

2 Literature Review

2.1 Random Matching with Limited Information

Rosenthal (1979) and Rosenthal and Landau (1979) introduced the study of repeated games with random matching. Rosenthal (1979) considered the special case of first-order information where players know only the action that their current opponent played in the previous period. He showed that Markovian equilibria exist, and that in the prisoner's dilemma cooperation can be supported by pure strategy equilibria only for a particular knife-edge value of the discount factor.[6]

Kandori (1992) and Ellison (1994) showed that cooperation in the prisoner's dilemma can be enforced by "contagion equilibria" in finite populations even when the players have "zero-order" information—that is, no information at all—about each other's past, but the required discount factor converges to 1 as the population becomes infinitely large. Kandori constructed simple contagion equilibria that exist only under a fairly strong restriction on the payoff functions—the loss parameter l needs to be sufficiently large, and in particular must diverge to ∞ as $\delta \to 1$. The issue is that if a player is very patient then even when he sees another player defect, and so knows that contagion has begun, he may still choose to cooperate to slow the spread of the contagion.[7] Ellison extended Kandori's results both to arbitrary payoff parameters and to approximately efficient equilibria in settings with a small amount of noise by using either public randomizing devices or "threading," which both serve to lower the players' effective discount factor, and so make it incentive-compatible for players to to carry out the punishments prescribed by the equilibrium. We find these sorts of strategies unintuitive, and some of the details of our model are designed to rule out them out.[8]

Three previous papers have studied cooperation in continuum-player repeated games with anonymous random matching and first-order information: Takahashi (2010), Heller and Mohlin (2018), and Bhaskar and Thomas (2018). Takahashi (2010) shows how cooperation can be supported when players know the entire record of each part-

[6]Rosenthal and Landau (1979) study two particular sorts of record-keeping technologies in the setting of an asymmetric battle of the sexes game. Their second, simpler, model has an index that goes up or down by 1 depending on which of the two actions the player uses, so it is related to the sorts of records we analyze in this paper; their first model of "comparative records" allows records to encode more than first-order information.

[7]Kandori also shows that cooperation can always be supported if higher-order information is available. Okuno-Fujiwara and Postlewaite (1995) derived a similar result in a continuum-population model. See also the literature on "standing" in evolutionary biology, starting with Sugden (1986) and more recently surveyed by Sigmund (2012).

[8]Deb, Sugaya, and Wolitzky (2018) 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.

ner's past play—all of the "first-order" information—but no higher-order information, such as the action that today's partner's partner took yesterday. Takahashi's general construction requires the use of somewhat complex and unintuitive "belief-free" mixed strategies. These strategies track the expected payoff of each player given her history of play and the equilibrium distribution of states and actions, and reduce a player's continuation payoff each time she plays D by just enough that she is always indifferent between C and D at every history. Because of this indifference, each player is willing to play C with exactly the probability corresponding to the current partner's record.

Takahashi also shows that first-order information allows cooperation to be supported in strict equilibria when the players are patient and the game is strictly supermodular. To do this, he follows Ellison and uses threads to calibrate the effective discount factor to be within the interval where players want to cooperate against a player who cooperates and defect against the player who defects. Conversely, Takahashi shows that no strategies support cooperation as a strict equilibrium when the payoffs are strictly submodular. The intuition is simple: with only first-order information, a player's continuation payoff depends on her record and what she does today but not on her current partner's action. Thus, in order for a player to strictly prefer to cooperate with a partner who will cooperate while defecting against a player who defects, payoffs must be strictly supermodular, and moreover the effective discount factor must be low enough that the difference in payoffs caused by the opponent's current action can offset the difference in future continuation payoffs.

Heller and Mohlin (2018)'s study of the prisoner's dilemma with anonymous random matching assumes that a small fraction of players are commitment types—which rules out belief-free equilibria—and that players live forever and are infinitely patient, but see only a finite sample of their partners' past actions. Players are restricted to use stationary strategies that condition only on the sampled actions of their partners and not on their own histories. Here, as in Takahashi, there are no cooperative equilibria when payoffs are submodular, and for much the same reason: players will play C or D depending on how likely it is their opponent plays C, so when their observation of their partner's play consists only of D's they are most likely to play C. But this means there is no incentive to play C, and cooperation breaks down. Conversely, Heller and Mohlin show that when payoffs are supermodular, the presences of commitment types allows the construction of an efficient and relatively simple mixed-strategy equilibrium without threads or public randomization, and that this equilibrium is essentially the only one that supports cooperation in their model.[9]

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

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

There is also a literature on repeated games with overlapping generations of non-anonynous players, e.g. Cremer (1986), Kandori (1992), Salant (1991), and Smith (1992). The lack of anonymity makes these papers less directly relevant.

[9]Heller and Mohlin also consider alternative information structures where players observe, for example a finite sample of their partners' past action *profiles*. The conditions for cooperation in this setting depend on payoff parameters via "snowballing" considerations similar to those in some of our results. However, the many differences between our models make the results difficult to compare directly. In another model with commitment types, Dilmé (2016) constructs a belief-free cooperative equilibrium for the case where $g = l$.

2.2 Simple Strategies in Fixed-Pair Interactions

Rubinstein (1986) and Abreu and Rubinstein (1988) introduced the study of repeated games played by automata. They assumed perfect monitoring, and as they acknowledge, their results are not robust to even a slight amount of noise. Compte and Postlewaite (2015) study fixed pairs playing the gift-exchange version of the prisoner's dilemma (which is on the boundary between the supermodular and submodular cases described above) with imperfect binary private signals of the partner's action. To model simplicity, they assume that strategies can be represented by automata with only two states, and they determine how the extent of cooperation is limited by the accuracy of the monitoring technology. Joe et al. (2012) computationally study whether "k-period mutual punishment strategies", which are conceptually similar to perfect tit-for-tat, can support cooperation in a two-player repeated prisoner's dilemma with private monitoring, and show that a simple version of this strategy can sustain fairly high payoffs for a range of parameters. Möbius (2001) and Olszewski and Safronov (2018a,b) consider simple "chips strategies" in repeated games, where a player gives her partner a token whenever he does her a favor, and a player who runs out tokens stops receiving favors until she reciprocates. These strategies are sometimes approximately efficient when players are patient.

Most evolutionary models of repeated games restrict to simple strategies to make the analysis tractable. This is the case for example in Axelrod and Hamilton (1981), who used evolutionary stability to argue that people will use the strategy "Tit-for-Tat" in the repeated prisoner's dilemma.[10]

Finally, the modern literature on repeated games in the laboratory provides evidence that simple strategies are used in play of fixed-partner repeated games (see e.g. Dal Bó and Fréchette (2018)). This helps motivate our interest in simple strategies.

[10] Axelrod (1984) showed that Tit-for-Tat and other simple strategies such as Tit-for-Two-Tats were selected in tournaments where participants submitted fully complete strategies to play the repeated game, but the requirement to submit strategies, as opposed to playing the extensive form version of the game period-by-period, may have led to simpler strategies than would otherwise have been used.

3 Steady-State Equilibria in Repeated Games with First-Order Information

Although most of the paper studies the prisoner's dilemma, we first present a model of first-order information in general stage games that may be of use in future work. We consider a discrete time model with a constant unit mass of players, each of whom has a geometrically-distributed lifespan with continuation probability $\gamma \in (0, 1)$, with exits balanced by a steady flow of new entrants. To motivate our exclusion of strategies that condition on calendar time, we assume the time horizon is doubly infinite.

Fix a finite symmetric game with action space A and payoff function $u : A \times A \to \mathbb{R}$. When players match they observe each other's record, which for now is an arbitrary integer $n \in \mathbb{Z}$. New players all enter with the same initial record n_0, which we set equal to 0 without loss of generality. One key restriction that we maintain throughout the paper is that these records track only first-order information—that is, data about how the player played—and do not depend on the play or records of the players she has been matched with.[11] The *state* of the system is then the share of players with each possible record; we denote this by $\mu \in \Delta(\mathbb{Z}) := M$.

To operationalize anonymous random matching in a continuum population, we specify that, when the current state of the system is μ, the distribution of matches is given by $\mu \times \mu$, so that, for each ordered pair $(n_1, n_2) \in \mathbb{Z}^2$, $\mu_{n_1}\mu_{n_2}$ is the fraction of matches between player 1's with record n_1 and player 2's with record n_2.

Definition 1. *A* record-keeping system *is a function $r : \mathbb{Z} \times A \to \Delta(\mathbb{Z})$ that specifies a probability distribution over a player's record tomorrow given the player's current record and her realized action in their current match. A record-keeping system r has* bounded-support updates *if there exists $B \in \mathbb{Z}$ such that support$(r(z, a))$ has at most B elements for all $(z, a) \in \mathbb{Z} \times A$.*

[11] This contrasts with the "status levels" studied by Okuno-Fujiwara and Postlewaite (1995) and the "standing" models of Sugden (1986) and Kandori (1991).

Here the stochastic term can represent errors in record-keeping, but it could also correspond to imperfect implementation of the intended action.[12]

We assume that when players meet each sees the record of her current opponent. In principle, each player can condition her play on the entire sequence of outcomes and past opponent records that she has seen. However, since we work in a model with a continuum of players, only the player's current record and that of her current partner matter for the player's current payoff, and only the player's own record will matter in the future. For this reason, all strict equilibria are *record-dependent,* meaning that they condition only on the player's current record and the record of her current partner. We write a record-dependent pure strategy as a function $s : \mathbb{Z}\times\mathbb{Z} \to A$, with the convention that the first coordinate is the player's own record and the second coordinate is that of the partner. Since we will restrict attention to strict equilibria, we consider only pure record-dependent strategies. Moreover, since every strict equilibrium in a symmetric, continuum-population model is symmetric, we also assume all players use the same strategy.

Given a record-keeping system r and a strategy s, we can define an update map $f_{r,s} : M \to M$ as follows: First, let $\phi(k', k'') := \mu_{k'}\mu_{k''}$ denote the probability that a player with record k' meets a player with record k''. Next, let $\rho_{r,s}(k', k'') \in \Delta(\mathbb{Z})$ be the probability distribution over next-period records of a player with record k' who meets a player with record k'', when all players use strategy s. Thus, $\rho_{r,s}(k', k'')[k] = r(k', s(k', k''))[k]$. Then, for $k \neq 0$, $f_{r,s}(\mu)[k] := \gamma \sum_{k',k''} \phi(k', k'')\rho_{r,s}(k', k'')[k]$, and $f_{r,s}(\mu)[0] := 1 - \gamma + \gamma \sum_{k',k''} \phi(k', k'')\rho_{r,s}(k', k'')[0]$. A *steady-state* under r, s is a state μ such that $f_{r,s}(\mu) = \mu$.

Theorem 1. *Under any record-keeping system with bounded-support updates and any record-dependent strategy, a steady state exists.*

The proof is in A.1 of the Appendix; all other omitted proofs can be found in either

[12] To restrict records to first-order information, we do not allow the update of a player's record to depend on the action that her partner played. In the implementation-errors interpretation of stochastic records, this requires that the signals have a product structure in the sense of Fudenberg, Levine, and Maskin (1994).

the Appendix (A) or the Online Appendix (OA). In outline, we 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 a player of age t, and let $\bar{M} = \left\{\mu \in M : \sum_{k \in R(t)} \mu_k \leq \gamma^t \ \forall t \in \mathbb{N}\right\}$. We first show $\bar{M}$ is compact in the sup norm: intuitively, bounded-support updates and geometrically distributed lifetimes imply that most players have records in the finite set $\bigcup_{t \leq T} R(t)$ for bounded T, so $\bar{M}$ resembles a finite-dimensional space. We then show that f maps $\bar{M}$ to itself and is continuous in the sup norm and note that $\bar{M}$ is convex, so we can appeal to a fixed point theorem.

Throughout the paper, our focus will be on equilibrium steady states. But note that Theorem 1 does not assert that the steady state for a given strategy is unique, and indeed we will see examples where it is not, as is the case under the strategy $Grim2$. Intuitively, this multiplicity corresponds to different initial conditions at time $t = -\infty$.

It remains to define equilibrium. Given a record-keeping system r, strategy s, and state μ, define the flow payoff of a player with record k as

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

Next, denote the probability that a player with record k today has record k' t periods from now by $\rho^t_{r,s,\mu}(k)[k']$: this is defined recursively by

$$\rho^1_{r,s,\mu}(k)[k'] = \sum_{k''} \mu_{k''} \rho_{r,s,\mu}(k,k'')[k']$$

and, for $t > 1$,

$$\rho^t_{r,s,\mu}(k)[k'] = \sum_{k''} \rho^{t-1}_{r,s,\mu}(k)[k''] \rho^1_{r,s,\mu}(k'')[k'].$$

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

$$V_{k,r,s,\mu} = (1-\gamma) \sum_{t=0}^{\infty} \gamma^t \sum_{k'} \rho^t_{r,s,\mu}(k)[k'] \pi_{k',r,s,\mu}.$$

Note that we have normalized continuation payoffs by $(1-\gamma)$ to express them in per-period terms. A pair (s,μ) is an *equilibrium* if μ is a steady-state under r,s and, for each own record k and opponent's record k', we have

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

In particular, a player's objective is to maximize her expected undiscounted lifetime payoff. An equilibrium is *strict* if the argmax is unique for all pairs of records (k,k'). Equilibrium existence follows immediately from Theorem 1.

Corollary 1. *Under any record-keeping system with bounded-support updates, an equilibrium exists.*

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

In contrast, the existence of a strict equilibrium is not guaranteed. A sufficient condition for strict equilibrium existence is for the stage game to have a strict and symmetric Nash equilibrium, as is the case with the prisoner's dilemma.

Corollary 2. *Under any record-keeping system with bounded-support updates, a strict equilibrium exists if the stage game has a strict and symmetric Nash equilibrium.*

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

4 Non-Trivial Equilibria in the Prisoner's Dilemma

For the rest of the paper we specialize to the case where the stage game is the prisoner's dilemma, given by Figure 1.[13] Note that *Always Defect* (i.e., $s(k,k')=D$ for all k,k') is

[13]The normalization in the figure is without loss: given a symmetric prisoner dilemma with payoff matrix $\begin{array}{ccc} & C & D \\ C & R,R & S,T \\ D & T,S & P,P \end{array}$, we can subtract P from all entries and divide by $R-P$, so that $g=$

a strict equilibrium for any parameter values. We call an equilibrium *non-trivial* if C is played with positive probability (i.e., $s(k, k') = C$ for some k, k' with $\mu_k, \mu_{k'} > 0$). This section shows that, as in Takahashi (2010) and Heller and Mohlin (2018), a non-trivial strict equilibrium can exist only if payoffs are strictly supermodular. This necessary condition applies for any first-order information structure.

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

Lemma 1. *At any record k in any strict equilibrium:*

1. *If $g = l$ then either the player is*
   - *an* unconditional defector *who plays D regardless of the opponent's record*
   - *an* unconditional cooperator *who plays C regardless of the opponent's record.*
2. *If $g > l$ then either the player is*
   - *an* unconditional defector *who plays D regardless of the opponent's record*
   - *an* unconditional cooperator *who plays C regardless of the opponent's record*

$(T-R)/(R-P)$ and $l = (P-S)/(R-P)$. If the stochastic nature of record updating comes from a probability ε that attempts to play C result in D, then the expected payoff matrix is slightly different, for example the payoff to (C, C) is $(1-\varepsilon)^2 R + \varepsilon(1-\varepsilon)(T+S) + \varepsilon^2 P$. In this case we denote the corresponding normalized payoffs by $g(\varepsilon)$ and $l(\varepsilon)$ respectively, where $g(0) = g$ and $l(0) = l$.

- *an* anti-reciprocator *who plays C against opponents who play D and plays D against opponents who play C.*

3. If $g < l$ then either the player is

- *an* unconditional defector *who plays D regardless of the opponent's record*
- *an* unconditional cooperator *who plays C regardless of the opponent's record*
- *a* reciprocator *who plays C against opponents who play C and plays D against opponents who play D.*

Proof. Fix a strict equilibrium. Because the records use only first-order information, each player's continuation payoff depends only on her current record and this period's play, so the optimal action in each period depends only on the player's record and the action prescribed by her opponent's record. There are four forms such a strategy could take; the lemma shows how which forms are possible in equilibrium depends on how g compares to l.

1. If $g = l$, the current period gain from playing D instead of C is independent of the opponent's action, so C is the strict best response against C iff it is also the strict best response to D, in which case the player is an unconditional cooperator. Similarly, D is the best response to C iff the player is an unconditional defector.

2. If $g > l$ and C is the strict best response to an opponent playing C, then playing C is also strictly optimal against an opponent playing D. Thus, there are no reciprocators, so every player is either an anti-reciprocator, an unconditional cooperator, or an unconditional defector.

3. If $g < l$ and D is strictly optimal against C, then D is also strictly optimal against D, so there are no anti-reciprocators. ■

Theorem 2. *If the payoffs are not strictly supermodular, then the unique strict equilibrium is* Always Defect*.*

Proof. Suppose that $g \geq l$. Fix a number k, and suppose two players who both have record k meet each other. By symmetry, they play either (C, C) or (D, D). In the former case, C is the strict best response to C, and therefore is also the strict best response to D. In the latter case, D is the strict best response to D, and therefore is also the strict best response to C. Thus there are no anti-reciprocators, and each player with record k is either an unconditional cooperator or an unconditional defector. Since this holds for all k, the distribution of opposing actions faced by any player is independent of her record. This implies D is always optimal. ■

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

5 Simple Records and Robust Equilibria

5.1 Counting D's

For the rest of the paper, we restrict attention to the following record-keeping system, which we call "counting D's": Newborn players have record 0. Whenever a player plays D, her record increases by 1. Whenever a player plays C, her record remains constant with probability $1-\varepsilon$, and her record increases by 1 with probability ε. Thus, a player's record is simply a count of the number of times she has been recorded as having played D.[14]

One motivation for considering this information structure is that counting D's is a simple record-keeping system that can support cooperation. Moreover, the equally simple system that counts C's instead of D's does not allow equilibria with any cooperation at all, as we show in the appendix. Roughly speaking, this is because a player with a very high C count would always play D, which by supermodularity implies that

[14]One could also consider a version of this information system with "two-sided errors," where a player's record remains constant with some probability $\varepsilon' > 0$ when she plays D. This would complicate notation without significantly affecting our results.

she must receive a low payoff, but then there is no incentive to acquire C's. Information structures that count both C's and D's can support cooperation, but are more complicated than counting only D's.

More importantly, we think counting D's is a natural information structure when players can forge records of past interactions but cannot hide true records. In particular, if players can create records of fake interactions where they played C, there cannot be an equilibrium where players are rewarded for having played C in the past. Under the plausible assumption that players are also not punished for having played C in the past, equilibrium strategies cannot depend on the number of times a player has played C at all. Thus, we consider record-keeping systems that only count D's.

Even with this minimalist record-keeping system, there can be equilibria that we regard as implausible and unintuitive. For this reason the next two subsections develop refinements on the equilibrium strategies. Here is an example of one type of equilibrium that we wish to rule out: Players with even records (including new players with record 0) are reciprocators and play C if and only if their partner's record is also even. Players with odd records are unconditional defectors and play D against all partners.

We will soon show that a family of strategies that generalize this "even-odd" form can sometimes be used to achieve efficient cooperation in the sense of the following definition.

Definition 2. **Limit efficiency** *is attainable if, for every* $\eta > 0$*, there exists* $\bar{\varepsilon} < 1$ *and a function* $\bar{\gamma} : (0,1) \to (0,1)$ *such that, whenever* $\varepsilon < \bar{\varepsilon}$ *and* $\gamma > \bar{\gamma}(\varepsilon)$*, there exists an equilibrium with* $V_0 > 1 - \eta$*.*

Note that V_0, the per-period expected payoff of a newborn agent, is also equal to average payoff in the population in every period. This follows because the expected fraction of a player's lifetime spent at record k is equal to the fraction of the population with record k (and there is no discounting, so both V_0 and the population-average payoff are given by $\sum_k \mu_k \pi_k$). Thus, when limit efficiency is attainable, the total population payoff approaches its maximum possible value in the iterated limit where first $\gamma \to 1$

and then $\varepsilon \to 0$.

Proposition 1. *Limit efficiency is attainable in strict equilibrium whenever the prisoner's dilemma is mild ($g < 1$).*

The proof of this result is in OA.2. Note that it leaves open whether and how limit efficiency might be possible when the prisoner's dilemma is severe ($g \geq 1$).

The proof relies on strategies that generalize even-odd. In fact, even-odd strategies themselves support limit efficiency in the case where $g < 1 < l$. To see this, note that, if all players use even-odd and the noise in records is small, then most players will be reciprocators, because odd records tend to change quickly. Thus, players' values are higher at even records, and so regardless of the parameters the prescribed play at odd records is optimal, because it both maximizes short-run payoffs and leads to a transition to a record with a higher payoff. Players with even records face different incentives, depending on their partner's record. When the partner has an even record (and is thus expected to play C), the one-shot gain g from defecting is less than the expected next-period loss of approximately 1 (when $\gamma \approx 1$ and $\varepsilon \approx 0$; note that with probability close to 1 the player's record in two periods will be even again whether or not she deviates today), so playing C as prescribed is optimal. When instead the partner has an odd record (and thus plays D), the one-shot gain from defecting is now l, which is greater than the next-period loss of 1, so playing D is optimal.[15]

5.2 Forgery-Proofness

Motivated by the idea that players can forge records of past interactions, we restrict attention to equilibria that are robust to adding D's to one's record.[16]

[15]If $g < l < 1$ then a one-period punishment for playing D is "too harsh," in that reciprocators then prefer to cooperate against defectors as well as cooperators. In this case, the construction must be modified so that, "on average," each play of D leads to a number $\rho < 1$ of punishment periods, where $\rho \in (g, l)$. For example, if players are reciprocators with record 0 or 1 mod 3, and unconditional defectors with record 2 mod 3, the average punishment is $\rho = .5$

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

Definition 3. *An equilibrium is* **forgery-proof** *if* $V_k \geq V_{k'}$ *for every* $k \leq k'$.

For example, even-odd violates forgery-proofness because a player with an odd record can avoid punishment by forging one interaction where she plays D. Note that even-odd has the flavor of the strategy "perfect tit-for-tat," which says to cooperate today if yesterday both partners played C or both partners played D, and to play D if exactly one partner played C. Moreover, like even-odd here, perfect tit-for-tat can be a strict equilibrium in the prisoner's dilemma with fixed pairs and imperfectly observed actions.[17] Both even-odd and perfect tit-for-tat are "forgiving" in the sense that the punishment phases are of finite length. Forgiveness could also be achieved with "repenting" strategies, where it is understood that if one player defects against the other, the defector will cooperate for a few rounds while their partner defects until the books are balanced and both players resume reciprocation. Although forgiving strategies have appealing properties in settings where records cannot be manipulated, these strategies cannot support cooperation in environments where D's can be forged. However, we will see that other non-trivial equilibria can satisfy forgery-proofness.

Forgery-proofness has a natural "dual" condition, which arises if, after taking her action and seeing what updated record the system assigns her, a player can choose whether to "erase" this update and return to previous record. In the current context, where a player's record either remains constant or increases by 1, this is simply the opposite of forgery-proofness: an equilibrium is *erasure-proof* if $V_k \leq V_{k'}$ for every $k \leq k'$. Under erasure-proofness, playing D maximizes both the stage game payoff and the continuation payoff, so the only erasure-proof equilibrium is *Always Defect*. In fact, we show in OA.4 that erasure-proofness rules out cooperation with any first-order information system, not just Counting D's.[18]

[17] "Ordinary" tit-for-tat is not an equilibrium with fixed pairs, imperfectly observed actions, and generic payoffs.

[18] The ability to erase individual records is more powerful than the ability to re-enter the game under a pseudonym, thus completely clearing one's record. In the latter case, for some information structures cooperation is possible via strategies where new players must "establish a reputation" for a number of periods before anyone will cooperate with them. See Friedman and Resnick (2001) for an analysis of "pseudonym-proofness" in a random-matching model with perfect monitoring.

5.3 Coordination-Proofness

We now define an additional robustness property, which we call *coordination-proofness*. The idea is that the equilibrium should not use "miscoordination" within a match as a threat to support social cooperation, because such an equilibrium will fall apart if matched partners do manage to coordinate successfully.

To formalize this property, let $\hat{u}_k(a, a')$ be the augmented payoff of a player with record k at action profile (a, a'), given by

$$\hat{u}_k(a, a') = (1-\gamma)u(a, a') + \gamma \sum_{k'} r(k, a)[k']V_{k'}.$$

Given our restriction to strict equilibria, a player with record k is a reciprocator iff

$$\hat{u}_k(C, C) - \hat{u}_k(D, C) > 0 > \hat{u}_k(C, D) - \hat{u}_k(D, D).$$

In other words, the augmented stage game played between any two reciprocators is a coordination game, where both (C, C) and (D, D) are stage-game Nash equilibria. It is straightforward to see that (C, C) is always the Pareto dominant equilibrium of this augmented game: since the gain to defection is always positive, $\hat{u}_k(C, C) - \hat{u}_k(D, C) > 0$ implies $V_k > V_{k+1}$, so for any reciprocator the (C, C) equilibrium yields higher continuation payoffs as well as higher stage-game payoffs. This observation motivates the following definition:

Definition 4. *An equilibrium is* ***coordination-proof*** *if whenever two reciprocators match, they play* (C, C).[19]

We will restrict attention to coordination-proof equilibria for the rest of the paper. An example of a strategy profile that satisfies forgery-proofness but violates coordination-proofness is the following: When players with records k and k' match, they

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

play (C, C) if $k + k' < 10$, and they play (D, D) otherwise. This violates coordination-proofness because any player with record $k \leq 9$ is a reciprocator, but two players with record 9 play D when matched with each other. The spirit of coordination-proofness is that this equilibrium is "unreasonable" because, if two partners with record 9 could agree to play (C, C) instead of (D, D) (holding fixed everyone else's strategy), both would have a strict incentive to abide by the agreement.

Our first result for coordination-proof equilibria is a partial converse to Proposition 1.[20]

Proposition 2. *If the prisoner's dilemma is severe ($g \geq 1$) then* Always Defect *is the unique coordination-proof strict equilibrium such that there is no record at which a player plays C for all opposing records (i.e., there are no unconditional cooperators).*

The introduction gave an intuition for Proposition 2 based on the "snowballing" of D's. A more detailed argument shows that we must have $\mu^R \geq g$, where μ^R is the fraction of reciprocators in the population, in any non-trivial equilibrium without unconditional cooperators. Thus, no such equilibrium can exist when $g \geq 1$.

To see this, let D_k be the occupation measure of D's faced by a player with record k.[21] Without unconditional cooperators, only (C, C) and (D, D) are played on path, so if a player with record k is a conditional cooperator, D_k must increase by at least g whenever she plays D (otherwise, she would never play C). Note that, when she matches with a defector, she both faces D today and plays D today.[22] Therefore,

$$D_k \geq \left(1 - \mu^R\right)(1 + g). \tag{1}$$

Moreover, if a player with record k is an unconditional defector, she faces D today for sure, and hence $D_k \geq D_{k+1} \geq \ldots \geq D_{k'}$, where k' is the least integer greater than

[20] More generally, we have shown that, for *any* record-keeping system (not just Counting D's), *Always Defect* is the unique coordination-proof strict equilibrium without unconditional cooperators that is robust to forging C's. Under the Counting D's information structure, all equilibria are trivially robust to forging C's, since a player's C count is completely irrelevant for her future play and payoffs.

[21] This is the "expected discounted number of D's".

[22] Supporting an equilibrium with no defectors—that is, $\mu^R = 1$—is impossible whenever $\varepsilon > 0$.

k such that a player with record k' is a reciprocator (and if $k' = \infty$, then $D_k = D_{k+1} = \ldots = 1$). Therefore, Equation 1 also holds for unconditional defectors, and hence must hold for everyone.

This implies that the total share of (D, D) outcomes, which equals $1 - \left(\mu^R\right)^2$, is at least $\left(1 - \mu^R\right)(1 + g)$. Thus,

$$1 - \left(\mu^R\right)^2 \geq \left(1 - \mu^R\right)(1 + g).$$

This requires $\mu^R \geq g$ when $\mu^R \neq 1$.

As we will see in Section 6, cooperation in the severe prisoner's dilemma is nonetheless possible under strategies with unconditional cooperators, such as $GrimKL$.

Coordination-proofness implies that every reciprocator plays C when matched with another reciprocator or an unconditional cooperator, and plays D when matched with a defector, so all reciprocators play C against the same set of opposing records. Therefore, under coordination-proofness, a strategy profile is completed characterized by a description of which records are reciprocators, which are unconditional cooperators, and which are defectors (formally, a mapping $\sigma : \mathbb{Z}_+ \to \{R, UC, D\}$).

This characterization simplifies the equilibrium conditions as follows: Fix an equilibrium (σ, μ). Recall that μ_k denotes the population share with record k. For each class of records $z \in \{R, UC, D\}$, let $\mathbb{K}_z \subset \mathbb{N}$ be the set of records in class z: that is, $\mathbb{K}_z = \{k \in \mathbb{N} : \sigma(k) = z\}$. Then, let μ^z denote the share of the population in class z, given by $\mu^z = \sum_{k \in \mathbb{K}_z} \mu_k$. Finally, we will also use the term *cooperator* for all players who are either reciprocators or unconditional cooperators (i.e., anyone who is not a defector), and we denote the population share of cooperators by $\mu^C = \mu^R + \mu^{UC} = 1 - \mu^D$.

The following result establishes how each player's record stochastically evolves from one period to the next (assuming the player survives to the next period).

Lemma 2. *Consider a steady state with share of cooperators μ^C.*

1. *Before being matched in the current period, the probability that a reciprocator with*

record k has record k at the end of the period is $(1-\varepsilon)\mu^C$, and the probability that she has record $k+1$ is $1-(1-\varepsilon)\mu^C$.

2. *Before being matched in the current period, the probability that an unconditional cooperator with record k has record k at the end of the period is $1-\varepsilon$ and the probability that she has record $k+1$ is ε.*

3. *Before being matched in the current period, the probability that a defector with record k has record $k+1$ at the end of the period is 1.*

Proof. Consider a reciprocator with record k. With probability μ^C, she will face another reciprocator in the upcoming period. In this case, she will play C, and her record will remain the same with probability $1-\varepsilon$ and increase to $k+1$ with probability ε. With probability $1-\mu^C$, she will face a defector. In this case, she will play D, and her record will increase to $k+1$ with probability 1. Thus, the probability that the reciprocator's record remains the same is the probability she faces another reciprocator times the probability that her action of C is correctly recorded, which equals $(1-\varepsilon)\mu^C$, while the probability that her record increases to $k+1$ is the complementary probability of $1-(1-\varepsilon)\mu^C$.

An unconditional cooperator with record k plays C regardless of her opponent's record, which causes her record to remain the same with probability $1-\varepsilon$ and to increase to $k+1$ with probability ε.

A defector always plays D, which always increases her record by 1. ■

We now characterize the incentive constraints of coordination-proof equilibria. We use the following notation: for a player with record k, $(C|C)_k$ is the condition that C is the best response to C, $(D|D)_k$ is the condition that D is the best response to D, $(C|D)_k$ is the condition that C is the best response to D, and $(D|C)_k$ is the condition that D is the best response to C. Note that, if a player with record k is a reciprocator, $(C|C)_k$ and $(D|D)_k$ must both hold; if a player with record k is an unconditional cooperator, $(C|D)_k$ must hold (which also implies $(C|C)_k$, by supermodularity); and if a player with record k is a defector, $(D|C)_k$ must hold (which also implies $(D|D)_k$).

Lemma 3.

1. *The $(C|C)_k$ constraint is*

$$\gamma(1-\varepsilon)(V_k - V_{k+1}) > (1-\gamma)g.$$

   The $(D|C)_k$ constraint is the opposite inequality.

2. *The $(D|D)_k$ constraint is*

$$\gamma(1-\varepsilon)(V_k - V_{k+1}) < (1-\gamma)l.$$

   The $(C|D)_k$ constraint is the opposite inequality

Proof. Consider a player with record k. When she plays C, her expected continuation payoff is $(1-\varepsilon)V_k + \varepsilon V_{k+1}$, while it is instead V_{k+1} when she plays D. Thus, the $(C|C)_k$ constraint is

$$1-\gamma+\gamma(1-\varepsilon)V_k + \gamma\varepsilon V_{k+1} > (1-\gamma)(1+g) + \gamma V_{k+1},$$

which is equivalent to the expression in Part 1 of the result. Likewise, the $(D|D)_k$ constraint is

$$\gamma V_{k+1} > -(1-\gamma)l + \gamma(1-\varepsilon)V_k + \gamma\varepsilon V_{k+1},$$

which is equivalent to the expression in Part 2. Clearly, $(D|C)_k$ is the opposite of $(C|C)_k$, and $(C|D)_k$ is the opposite of $(D|D)_k$. ■

6 *GrimKL* Strategies

6.1 Combining Forgery-Proofness and Coordination-Proofness

The main result of this subsection characterizes the unique family of strategies that can be equilibria that satisfy both forgery-proofness and coordination-proofness. This is the family of *GrimKL* strategies, which is defined for non-negative integers K and L so that a player is a reciprocator for the first K records, $0 \le k \le K-1$, an unconditional cooperator for the next L records, $K \le k \le K+L-1$, and a defector for all other records, $k \ge K+L$. Note that *GrimKL* with $K = L = 0$ is the familiar *Always Defect* strategy, and *GrimKL* with $L = 0$ is *GrimK*: here a player is a reciprocator for the first K records, $0 \le k \le K-1$, and a defector for all other records, $k \ge K$, but never an unconditional cooperator.[23]

Figure 2 depicts the steady state record shares, μ_k, and the value functions, V_k, in a *GrimKL* equilibrium for $\gamma = .95$, $\varepsilon = .05$, $g = .5$, and $l = 1$; here $K = 4$ and $L = 1$. In this equilibrium, $\mu^R \approx .492$, $\mu^{UC} \approx .260$, and $\mu^D \approx .247$. We note that there is no non-trivial *GrimK* equilibrium for these parameters, which can be shown using the characterization of *GrimK* equilibria in A.6. Intuitively, the fact that players with $k = 4$ always play C implies that transitions from record 4 to record 5 are unlikely, which keeps the total share of cooperators high enough for a non-trivial *GrimKL* equilibrium to exist, whereas with *GrimK* strategies transitions to the defector region occur too quickly to support positive steady-state cooperation.

We now present the characterization result.

Theorem 3. *Any strict equilibrium that satisfies forgery-proofness and coordination-proofness corresponds to GrimKL for some K, L. Moreover, if for some K, L, GrimKL is a strict equilibrium, it satisfies forgery-proofness and coordination-proofness.*

The result can be understood in three steps. First, forgery-proofness implies that there is a critical record $\bar{k}$ such that all records $k \ge \bar{k}$ are defectors. Intuitively, this

[23]Note that *Always Defect* is a special case of *GrimK* as well as *GrimKL*.

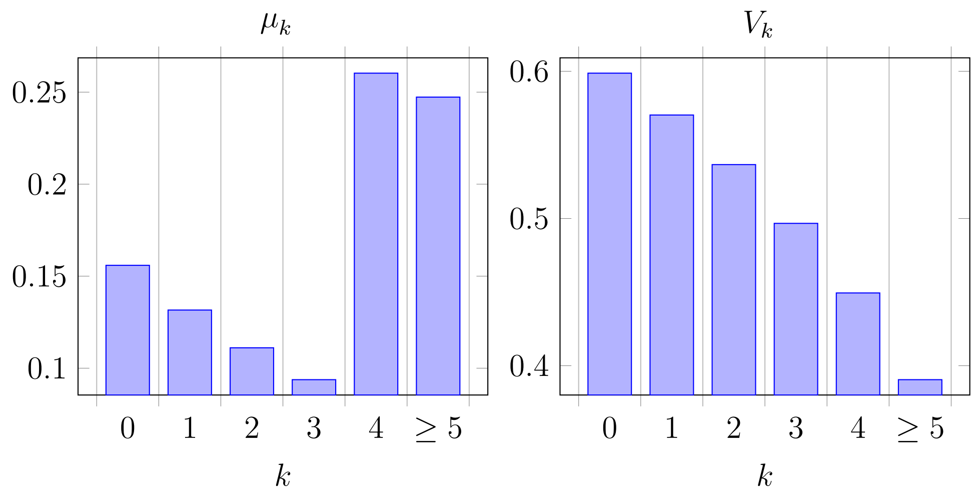

Figure 2: Steady state record shares and valuation functions for *Grim*41.

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

The following two subsections will analyze the maximum level of cooperation that can be supported by the family of $GrimKL$ strategies. We first define the maximal equilibrium level of cooperation for both the $GrimKL$ and $GrimK$ strategy families. Let $\overline{\mu}^C_{KL}(\gamma, \varepsilon)$ be the maximal share of cooperators that is attained in any $GrimKL$ equilibrium for parameters γ and ε:

$$\overline{\mu}^C_{KL}(\gamma, \varepsilon) = \sup\{\mu^C : \mu^C \text{ is the share of cooperators in a } GrimKL \text{ equilibrium}\}.$$

Likewise, let $\overline{\mu}^C_K(\gamma, \varepsilon)$ be the maximal share of cooperators that is attained in any $GrimK$ equilibrium for parameters γ and ε:

$$\overline{\mu}^C_K(\gamma, \varepsilon) = \sup\{\mu^C : \mu^C \text{ is the share of cooperators in a } GrimK \text{ equilibrium}\}.$$

6.2 $GrimK$ Strategies

We first focus on the case of $GrimK$, where unconditional cooperators are not present.

Theorem 4.

1. *For $g < 1$ and $l > g/(1-g)$,*

$$\lim_{\varepsilon \to 0} \lim_{\gamma \to 1} \overline{\mu}^C_K(\gamma, \varepsilon) = \frac{l}{1+l}.$$

2. *For $g < 1$ and $l \leq g/(1-g)$,*

$$\lim_{\varepsilon \to 0} \lim_{\gamma \to 1} \overline{\mu}^C_K(\gamma, \varepsilon) = 0.$$

3. *For $g \geq 1$, $\overline{\mu}^C_K(\gamma, \varepsilon) = 0$ for all $(\gamma, \varepsilon) \in (0,1) \times (0,1)$.*

Part 3 of Theorem 4 says that no cooperation is possible in a $GrimK$ equilibrium whenever $g \geq 1$, regardless of γ and ε. Thus, *Always Defect* (i.e. $Grim0$) is the unique equilibrium in this family in a severe prisoner's dilemma. This is a consequence of Proposition 1, since with *GrimK* there are no unconditional cooperators.

Part 1 of Theorem 4 says that in a mild prisoner's dilemma with $g < 1$ and $l > g/(1-g)$, *GrimK* strategies can support strictly positive cooperation, and moreover that the maximal level of cooperation tends to $l/(1+l)$ in the iterated limit where γ approaches 1 and then ε approaches 0. Part 2 says that for $g < 1$ and $l < g/(1-g)$ the level of cooperation approaches 0 in the iterated limit. The reason a higher l makes cooperation easier to sustain in a *GrimK* equilibrium is that it makes it it easier to satisfy the $(D|D)$ constraints for higher values of μ^C.[24]

To understand this result, note that the $(D|D)_{K-1}$ constraint may be written as

$$\mu^C < \frac{1}{\gamma(1-\varepsilon)} \frac{l}{1+l}. \tag{2}$$

This comes from combining the recursive equation for V_{K-1},

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

(which follows because, once a player with record $K-1$ gains another D in her record, she receives payoff 0 forever), with the $(D|D)_{K-1}$ constraint

$$0 > (1-\gamma)(-l) + \gamma(1-\varepsilon)V_{K-1}.$$

As $\gamma \to 1$ and $\varepsilon \to 0$, Equation 2 yields the upper bound $\mu^C < \frac{l}{1+l}$.

On the other hand, since *GrimK* does not involve unconditional cooperators, by the discussion following Proposition 2 we must have $\mu^C > g$. There exists such a value of μ^C that also satisfies the upper bound $\mu^C < \frac{l}{1+l}$ if and only if $l > \frac{g}{1-g}$.

We note that *Grim*2 provides a simple example of a strategy profile with multiple associated steady states. For example, if $\gamma = .8$ and $\varepsilon = .01$ then there are three steady-

[24]This is essentially the same reason that a high value of l is necessary for contagion strategies to form an equilibrium in Kandori (1992).

state values of μ^C under $Grim2$, given by $\mu^C \approx .6738$, $\mu^C \approx .8536$, and $\mu^C \approx .9979$. Additionally, for $g < .171$ and $l > 3.77$, $Grim2$ is an equilibrium for all three of these values of μ^C.

6.3 General $GrimKL$ Strategies

We now turn to the full family of $GrimKL$ strategies, where L can be greater than 0. Let $b : \mathbb{R}_+ \to \mathbb{R}$ be the function given by

$$b(g) = \min\left\{\frac{1+g}{|\ln(1+g)-1|}, 21.9223 - 3.57143g\right\}.^{25}$$

Theorem 5.

1. *For* $l > \max\{g(g+1), b(g)\}$, $\lim_{\varepsilon\to 0}\liminf_{\gamma\to 1}\overline{\mu}^C_{KL}(\gamma,\varepsilon) = 1$.

2. *For* $l \le g(g+1)$, $\lim_{\gamma\to 1}\overline{\mu}^C_{KL}(\gamma,\varepsilon) = 0$ *for all* $\varepsilon \in (0,1)$.

3. *For* $g \ge 1$ *and* $l \le g(g+1)$, $\overline{\mu}^C_{KL}(\gamma,\varepsilon) = 0$ *for all* $(\gamma,\varepsilon) \in (0,1)\times(0,1)$.

Remark 1: Interpreting the parametric conditions of Theorem 6. Part 3 of Theorem 5 says that, when $g \ge 1$ and $l \le g(g+1)$ so that the prisoner's dilemma is both severe and mildly supermodular, *Always Defect* is the unique equilibrium in the $GrimKL$ family. Part 2 says that when the prisoner's dilemma is mildly supermodular, the maximal level of cooperation approaches 0 as $\gamma \to 1$, for any value of ε, regardless of whether the prisoner's dilemma is severe.

In contrast, Part 1 of Theorem 5 says that when the prisoner's dilemma is sufficiently supermodular, the maximal level of cooperation with $GrimKL$ strategies tends to 1 in the iterated limit. In particular, for sufficiently severe prisoner's dilemmas with $g(g+1) > b(g)$ (i.e., $g >\approx 2.858$), limit efficiency is achieved if $l > g(g+1)$. For smaller values of g, we use the stronger requirement that $l > b(g)$ to guarantee limit efficiency. We do not know whether this condition is necessary for limit efficiency; it

[25] $21.9223 - 3.57143g$ is the approximate form of $(3e^{\phi} - 2 - 2g)/(\phi - 1)$ when $\phi = 1.56$.

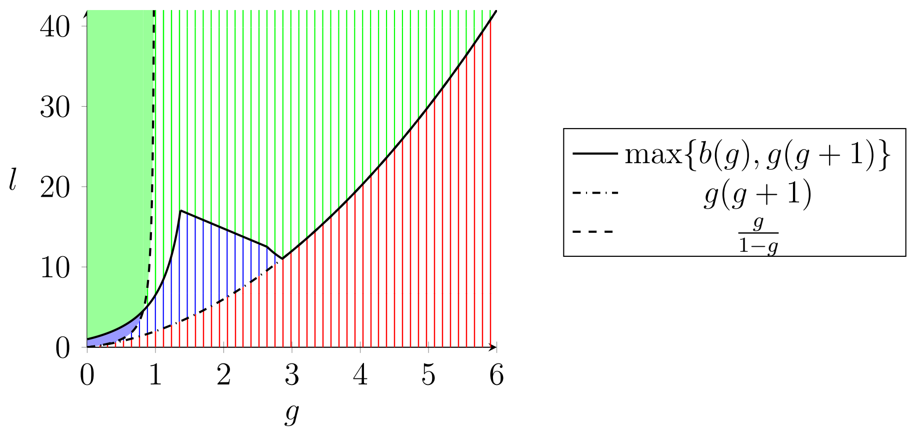

Figure 3: The performance of the maximally efficient $GrimKL$ and $GrimK$ strategies in the iterated limit where γ approaches 1 and then ε approaches 0. $GrimKL$ attains limit efficiency in the green regions and has limit cooperation of 0 in the red region. Our results do not pin down the limit performance of $GrimKL$ in the blue regions. Solid shading indicates that $GrimK$ attains strictly positive limit cooperation; vertical lines indicate that $GrimK$ has limit cooperation of 0.

comes from the fact that K and L must be integers, which affects the feasibility of generating particular values of μ^C, μ^R, μ^{UC}, as is discussed in A.7. Our results do not say much about cooperation in the iterated limit when $l > g(g+1)$ holds, but $l > b(g)$ does not. However, if integer problems could be ignored, then $l > g(g+1)$ would again suffice for limit efficiency. If this were the case, for any parameters such that the $GrimK$ family permits cooperation in the iterated limit, the $GrimKL$ family would enable cooperation in the iterated limit with strictly smaller values of l because $g(g+1) < g/(1-g)$ for all $g < 1$. Figure 3 presents a graphical illustration of the performance of $GrimKL$ and $GrimK$ in the iterated limit.

Remark 2: Why limit efficiency requires $l > g(g+1)$. The inequality $l > g(g+1)$ comes from combining the constraints $\mu^R > 1 - \frac{1}{1+g}$ and $\mu^R < 1 - \frac{g}{l}$. We now explain each of these in turn.

The condition $\mu^R > 1 - \frac{1}{1+g}$ is necessary for the flow payoff of an unconditional cooperator to exceed the flow payoff of a defector, which in turn is necessary for $GrimKL$ to be an equilibrium. (Otherwise, unconditional cooperators would prefer to defect.) In particular, unconditional cooperators earn a higher flow payoff than defectors if and only if $\mu^C - \left(1-\mu^C\right) l > \mu^{UC}(1+g)$, and this is equivalent to $\mu^R(1+l) + \mu^{UC}(l-g) > l$. The left-hand side of this inequality is increasing in μ^{UC}, so a necessary condition is that it holds when μ^{UC} takes on its highest possible value of $1-\mu^R$. Thus, substituting $\mu^{UC} = 1-\mu^R$, the inequality becomes $\mu^R(1+l) + \left(1-\mu^R\right)(l-g) > l$, which is equivalent to $\mu^R > 1 - \frac{1}{1+g}$.

The condition $\mu^R < 1 - \frac{g}{l}$ is necessary for $(C|C)_0$ and $(D|D)_{K-1}$ to be satisfied simultaneously. In particular, the probability that a player survives to reach record $K-1$ is the ratio of the benefit of playing C with record 0 to the benefit of playing C with record $K-1$, and it is also approximately equal to $1-\mu^R$, the share of players with record $\geq K$.[26] Thus,

$$\frac{\text{benefit of playing } C \text{ at } k=0}{\text{benefit of playing } C \text{ at } k=K-1} \approx 1-\mu^R.$$

On the other hand, we have

$$\frac{\text{cost of playing } C \text{ vs. } C}{\text{cost of playing } C \text{ vs. } D} = \frac{g}{l}.$$

Now, for it to be optimal both to play C vs. C with record $k=0$ and to play D vs. D with record $k=K-1$, it must be the case that

$$\frac{\text{benefit of playing } C \text{ at } k=0}{\text{cost of playing } C \text{ vs. } C} > \frac{\text{benefit of playing } C \text{ at } k=K-1}{\text{cost of playing } C \text{ vs. } D}.$$

Combining these observations, we have

$$1-\mu^R > \frac{g}{l} \iff \mu^R < 1-\frac{g}{l}.$$

Note also that the constraint $\mu^R < 1-\frac{g}{l}$ implies that, in any equilibrium that supports $\mu^C \approx 1$, we must have $\mu^{UC} > \frac{g}{l}$. Thus, unless l is much larger than g, there

[26] This follows because $1-\mu^R$ is the probability that a player survives to reach record K, and when γ is large relative to ε these two probabilities are almost the same.

must be a substantial share of unconditional cooperators in any efficient equilibrium.

Remark 3: The role of unconditional cooperators. Theorem 5 demonstrates the power of unconditional cooperation in sustaining equilibrium cooperation. Intuitively, unconditional cooperators help support cooperation in two ways. First, there is the relatively straightforward effect that turning those reciprocators who are most tempted to play C against D into unconditional cooperators relaxes the $(D|D)$ constraint. This explains why $GrimKL$ can sometimes support almost full cooperation, while $GrimK$ can never support $\mu^C > \frac{l}{1+l}$.

Second, there is the novel effect that the presence of unconditional cooperators increases the steady-state share of cooperators by slowing the transition from cooperator records to defector records, and in particular cuts off the snowballing D's that can destroy the possibility of positive steady-state cooperation in severe prisoner's dilemmas. We show in Proposition 4 in the Appendix that the the $(C|C)$ constraint cannot be satisfied if $g \geq 1$, so it is distinct from the $(D|D)$ constraint. For this reason, when $g \geq 1$ then $GrimK$ strategies cannot support positive cooperation *even in the relaxed problem where players' incentives when their opponents play D are completely ignored.*[27] Thus, this benefit of unconditional cooperators has nothing to do with relaxing the $(D|D)$ constraint.

Remark 4: Proving Theorem 5. The proof of Theorem 5 is more involved than that of Theorem 4. The main additional complication is that the cost of increasing one's record by 1 is very similar for a player with record $K-1$ and a player with record K, which makes it hard to find K and L that simultaneously satisfy the $(D|D)_{K-1}$ and $(C|D)_K$ constraints. This issue is also the source of the integer problems that lead to the requirement that $l > b(g)$ in Part 1 of Theorem 5. We provide a detailed proof outline at the beginning of A.7.

[27]Moreover, the proof of the more general Proposition 2 uses only the $(C|C)$ constraints and not the $(D|D)$ constraints.

7 Conclusion

We conclude with some observations about extensions and alternative models.

Other types of equilibria. With the classes of strategies we have discussed so far, it is necessary that either $g < 1$ or $l > g(g+1)$ for any cooperation to occur in equilibrium. This is not true in general if forgery-proofness and coordination-proofness are not imposed. Consider the family of strategies defined for four non-negative integers, K_1, K_2, K_3, and K_4, as follows: A player is a reciprocator for the first K_1 records, a defector for the next K_2 records, a reciprocator again for the next K_3 records, an unconditional cooperator for the next K_4 periods, and a defector for all higher records. We show in OA.11 that strategies in this family can sometimes enable cooperation in equilibrium even when $g > 1$ and $l < g(g+1)$. Our intuition for this result is that, by splitting the set of defector records into two classes, these strategies accelerate the punishment for acquiring D's and thus strengthen incentives for the first class of reciprocators. However, players in the first defector class would gain by forging D's.

Endogenous Stochastic transitions. We have restricted attention to equilibria with exogeneous record transitions. Stochastic transitions that can be tuned to the game parameters can be used as randomization devices, and improve on $GrimK$ and $GrimKL$ strategies in at least two ways. First, a stochastic version of $GrimK$ can attain full limit efficiency when $g < 1$: the intuition is that making transitions from the reciprocator state to the defector state stochastic relaxes the $(D|D)$ constraint, which was responsible for the $\mu^C \leq \frac{l}{1+l}$ upper bound for deterministic strategies, as well as the requirement that $l > g/(1-g)$.[28] Second, a stochastic version of $GrimKL$ can avoid the integer problems that arose in Theorem 6, and thus relax the condition for limit efficiency from $l > \max\{g(g+1), b(g)\}$ to $l > g(g+1)$.[29] We demonstrate these results

[28] We prove this with a strategy with only two records, where record-0 players are reciprocators and record-1 players are defectors.

[29] We prove this with a strategy with three records, where record-0 players are reciprocators, record-1 players are unconditional cooperators, and record-2 players are defectors. As with $GrimKL$, the defectors remain defectors forever; the difference is that playing D causes one's record to increase with probability less than 1. We do not know whether more complex stochastic transition rules could allow limit efficiency for a wider range of parameters.

in OA.12, using transition probabilities between records that depend on the parameters ε and γ.[30]

Higher-order information. We have also restricted to record-keeping systems that use only first-order information. If the update of a player's record can depend on her opponent's action as well as her own, supporting cooperation becomes much easier: for example, if a player's record increases only when she plays D versus C, this cuts off the "snowballing" problem, and we prove in OA.13 that limit efficiency can be attained by $GrimK$ strategies for any values of g and l, even those with $g > l$. This shows that the key limitation of first-order information is that, when a player sees that her opponent has played a certain number of D's, she cannot tell if these plays correspond to "selfish" deviations or "deserved" punishments.[31]

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

Different stage games. Our model of random-matching with first-order information could also be used to study stage games other than the prisoner's dilemma. It is also easy to adapt our model to settings with multiple populations of players. For example, suppose a population of player 1's and a population of player 2's repeatedly play the product choice game, where only player 1 faces binding moral hazard at the efficient action profile (and player 2 wants to match player 1's action). In this game, efficient

[30] These results consider the problem of designing an optimal record-keeping system within some class, while the rest of this paper studies a particular simple and natural record-keeping system. See Bhaskar and Thomas (2018) for an information-design perspective in a lending game model.

[31] In a related result, Heller and Mohlin (2018) show that, in their model, full cooperation is always possible when players observe only their partners' past actions when facing C.

payoffs can always be supported in the limit where first $\gamma \to 1$ and then $\varepsilon \to 0$, using the following adaptation of $GrimK$ strategies: in each match, both partners play C if player 1's record is $k < K$, and both play D if player 1's record is $k \geq K$. This example suggests our model might be especially tractable in settings with 1-sided moral hazard.

References

[1] Dilip Abreu and Ariel Rubinstein. "The Structure of Nash Equilibrium in Repeated Games with Finite Automata". *Econometrica* (1988).

[2] Charalambos D. Aliprantis and Kim C. Border. *Infinite Dimensional Analysis: A Hitchhiker's Guide 3rd Edition*. Third. Berlin: Springer, 2006.

[3] Robert M. Axelrod. *The evolution of cooperation*. Basic Books, 1984.

[4] Robert Axelrod and William D. Hamilton. "The evolution of cooperation." *Science (New York, N.Y.)* (1981).

[5] B Douglas Bernheim and Debraj Ray. "Collective Dynamic Consistency in Repeated Games". *Games and Economic Behavior* (1989).

[6] V Bhaskar and Caroline Thomas. "Community Enforcement of Trust with Bounded Memory". *Review of Economic Studies* (2018).

[7] Sylvain Chassang and Satoru Takahashi. "Robustness to incomplete information in repeated games". *Theoretical Economics* (2011).

[8] Judith A. Chevalier and Dina Mayzlin. "The Effect of Word of Mouth on Sales: Online Book Reviews". *Journal of Marketing Research* (2006).

[9] Olivier Compte and Andrew Postlewaite. "Plausible cooperation". *Games and Economic Behavior* (2015).

[10] Jacques Cremer. "Cooperation in Ongoing Organizations". *The Quarterly Journal of Economics* (1986).

[11] Pedro Dal Bó and Guillaume R Fréchette. "On the Determinants of Cooperation in Infinitely Repeated Games: A Survey". *Journal of Economic Literature* (2018).

[12] Joyee Deb, Takuo Sugaya, and Alexander Wolitzky. "The Folk Theorem in Repeated Games with Anonymous Random Matching". 2018.

[13] Francesc Dilmé. "Helping Behavior in Large Societies". *International Economic Review* (2016).

[14] Robert C. Ellickson. *Order without Law : How Neighbors Settle Disputes*. Harvard University Press, 1994.

[15] Glenn Ellison. "Cooperation in the prisoner's dilemma with anonymous random matching". *Review of Economic Studies* (1994).

[16] Joseph Farrell and Eric Maskin. "Renegotiation in Repeated Games". *Games and Economic Behavior* (1989).

[17] Eric J. Friedman and Paul Resnick. "The Social Cost of Cheap Pseudonyms". *Journal of Economics and Management Strategy* (2001).

[18] Drew Fudenberg, David Levine, and Eric Maskin. "The Folk Theorem with Imperfect Public Information". *Econometrica* (1994).

[19] Drew Fudenberg and Eric Maskin. "Evolution and Cooperation in Noisy Repeated Games". *The American Economic Review* (1990).

[20] Drew Fudenberg, David G. Rand, and Anna Dreber. "Slow to Anger and Fast to Forgive: Cooperation in an Uncertain World". *American Economic Review* (2012).

[21] Yuval Heller and Erik Mohlin. "Observations on Cooperation". *Review of Economic Studies* (2018).

[22] Yongjoon Joe et al. "Automated Equilibrium Analysis of Repeated Games with Private Monitoring: A POMDP Approach". *Conference on Autonomous Agents and Multiagent Systems*. 2012.

[23] Michihiro Kandori. "Social Norms and Community Enforcement". *The Review of Economic Studies* (1992).

[24] Michael Luca and Georgios Zervas. "Fake It Till You Make It: Reputation, Competition, and Yelp Review Fraud". *Management Science* (2016).

[25] Dina Mayzlin, Yaniv Dover, and Judith Chevalier. "Promotional Reviews: An Empirical Investigation of Online Review Manipulation". *American Economic Review* (2014).

[26] Paul Milgrom, Douglass C. North, and Barry R. Weingast. "The Role of Institutions in the Revival of Trade: The Law Merchant, Private Judges, and the Champagne Fairs". *Economics and Politics* (1990).

[27] Markus M Möbius. "Trading Favors". 2001.

[28] Martin A. Nowak and Karl Sigmund. "Evolution of indirect reciprocity by image scoring". *Nature* (1998).

[29] Masahiro Okuno-Fujiwara and Andrew Postlewaite. "Social Norms and Random Mathcing Games". *Games and Economic Behavior* (1995).

[30] Wojciech Olszewski and Mikhail Safronov. "Efficient chip strategies in repeated games". *Theoretical Economics* (2018).

[31] Wojciech Olszewski and Mikhail Safronov. "Efficient cooperation by exchanging favors". *Theoretical Economics* (2018).

[32] Karthik Panchanathan and Robert Boyd. "A tale of two defectors: the importance of standing for evolution of indirect reciprocity". *Journal of Theoretical Biology* (2003).

[33] Robert W. Rosenthal. "Sequences of Games with Varying Opponents". *Econometrica* (1979).

[34] Robert W. Rosenthal and Henry J. Landau. “A game-theoretic analysis of bargaining with reputations”. *Journal of Mathematical Psychology* (1979).

[35] Ariel Rubinstein. “Finite automata play the repeated prisoner’s dilemma”. *Journal of Economic Theory* (1986).

[36] David J. Salant. “A repeated game with finitely lived overlapping generations of players”. *Games and Economic Behavior* (1991).

[37] Karl Sigmund. “Moral assessment in indirect reciprocity.” *Journal of theoretical biology* (2012).

[38] Lones Smith. “Folk Theorems in Overlapping Generations Games”. *Games and Economic Behavior* (1992).

[39] Robert Sugden. *The Economics of Rights, Co-operation and Welfare*. Oxford: Basil Blackwell, 1986.

[40] Satoru Takahashi. “Community enforcement when players observe partners’ past play”. *Journal of Economic Theory* (2010).

Appendix

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

A.1 Proof of Theorem 1

Let B be a bound on the update supports. Without loss, relabel records so that two players with different ages can never share the same record. Then let $R(t)$ be the set of feasible records for a player of age t, and let $\bar{M} = \left\{\mu \in M : \sum_{k \in R(t)} \mu_k \leq \gamma^t \ \forall t \in \mathbb{N}\right\}$. Clearly, $\bar{M}$ is convex. Section OA.1 establishes the following three claims:

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

Claim 2. *$f_{r,s}$ maps $\bar{M}$ to itself.*

Claim 3. *$f_{r,s}$ is continuous in the* sup *norm topology.*

Since every normed space is a locally convex Hausdorff space, the theorem follows from Corollary 17.56 (page 583) of Aliprantis and Border (2006). The key point is that the combination of bounded-support updates and geometric player lifetimes means that the tails of the feasible population states must be thin. This explains why $\bar{M}$ is compact, and also plays a key role in ensuring that f is continuous.

A.2 Impossibility for Counting C's

We show cooperation cannot be supported when a player's record is the number of times she has played C (where a play of C is recorded as D with probability ε). We continue to assume $g < l$, as otherwise Theorem 2 already implies that cooperation is impossible.

Proposition 3. *When records count C's, the unique strict equilibrium is* Always Defect*.*

Proof. Note that, in any strict equilibrium, if records k and $k+1$ are both reached with positive probability, then $V_k \leq V_{k+1}$. For if $V_k > V_{k+1}$, then a player with record k always plays D (since this maximizes both her flow payoff and her continuation payoff), but since plays of D are always recorded accurately this implies no player ever reaches record $k+1$.

We now show that there exists K^* such that a player with record $k > K^*$ always defects. To see this, let $V^* = \sup_k V_k$. For a player with record k, the gain in continuation value from playing C is is at most $V^* - V_k$, while the gain in flow utility from playing D is at least $(1-\gamma)\,g$. Hence, whenever $V_k > V^* - (1-\gamma)\,g$, a player with record k always defects. Since V_k is non-decreasing, there exists K^* such that a player with record $k > K^*$ always defects.

By supermodularity, everyone other than unconditional cooperators play D against a player with record $k > K^*$. Therefore, $V_{K^*+1} = V^*$, and this is precisely the payoff a player receives from always playing D while facing C from unconditional cooperators and D from everyone else. Note that this is also a lower bound on the payoff attained by a player who always plays D from the beginning of the game. Therefore, $V_0 \geq V^*$, which since V_k is non-decreasing implies $V_0 = V^*$. But this implies that V_k is constant, so playing D is always optimal. ■

The above argument uses the fact that plays of D are perfectly recorded to establish that V_k is non-decreasing. However, as we show in Section OA.3, the same conclusion applies even if playing D causes the record to increase with probability $1 - \varepsilon' < 1$.

A.3 Proof of Proposition 2

Let $\mathbb{K}_R$ denote the set of reciprocator records, and let $\mu^R = \sum_{k \in \mathbb{K}_R} \mu_k$ be the share of reciprocators. Since only (C, C) and (D, D) are played on path, coordination-proofness implies that the flow payoff of a player with record k is μ^R if $k \in \mathbb{K}_R$ and 0 if $k \notin \mathbb{K}_R$. Thus,

$$V_0 = \sum_{k \in \mathbb{N}} \mu_k \pi_k = \sum_{k \in \mathbb{K}_R} \mu_k \mu^R = \left(\mu^R\right)^2. \tag{3}$$

play (C, C).

Suppose a player with record k is a reciprocator: $k \in \mathbb{K}_R$. The corresponding $(C|C)_k$ constraint gives

$$V_k - V_{k+1} > \frac{1}{1-\varepsilon} \frac{1-\gamma}{\gamma} g. \tag{4}$$

On the other hand,

$$V_k = (1-\gamma)\mu^R + \gamma(1-\varepsilon)\mu^R V_k + \gamma(1-(1-\varepsilon)\mu^R)V_{k+1}.$$

Grouping terms, we obtain

$$\begin{aligned}
(1-\gamma) V_k &= (1-\gamma)\mu^R - \gamma\left(1-(1-\varepsilon)\mu^R\right)(V_k - V_{k+1}) \\
&< (1-\gamma)\mu^R - \gamma\left(1-(1-\varepsilon)\mu^R\right)\left(\frac{1}{1-\varepsilon}\frac{1-\gamma}{\gamma}g\right) \quad \text{(by (4))} \\
&= (1-\gamma)\mu^R - (1-\gamma)\left(\frac{1}{1-\varepsilon} - \mu^R\right) g \\
&\Longleftrightarrow \\
V_k &< (1+g)\mu^R - \frac{g}{1-\varepsilon}. \qquad (5)
\end{aligned}$$

Next, we establish that either $\mathbb{K}_R = \emptyset$ (in which case the conclusion of the proposition is immediate) or

$$V_0 < (1+g)\mu^R - \frac{g}{1-\varepsilon}. \qquad (6)$$

If $0 \in \mathbb{K}_R$, (6) is simply (5) with $k = 0$. So suppose instead that $0 \notin \mathbb{K}_R$ and $\mathbb{K}_R \neq \emptyset$. Recall that $k \notin \mathbb{K}_R$ means that $\pi_k = 0$. Therefore, $k \notin \mathbb{K}_R$ implies that $V_k \leq V_{k+1}$. If we let k' be the smallest record such that $k' \in \mathbb{K}_R$ (which exists since $\mathbb{K}_R \neq \emptyset$), we have $V_0 \leq V_{k'}$, so (6) follows since (5) holds for $k = k'$.

Combining (3) and (6) gives

$$\begin{aligned}
\left(\mu^R\right)^2 &< (1+g)\mu^R - \frac{g}{1-\varepsilon} \\
&\Longleftrightarrow \\
-\left(\mu^R\right)^2 + (1+g)\mu^R - \frac{g}{1-\varepsilon} &> 0.
\end{aligned}$$

If $g \geq 1$ the left-hand side of this inequality is increasing in μ^R, yet at $\mu^R = 1$ it is negative. Thus, this inequality cannot be satisfied for any $\mu^R \in [0,1]$ when $g \geq 1$.

A.4 Lemmas about Coordination-Proof Equilibria

Define i_k to be the total inflow of players into record k. This equals $1-\gamma$ for $k=0$ since all newborn players have record $k=0$, while i_k equals the share of players that move into record k from record $k-1$ in a period for $k>0$. Define τ_k to be the total share of players that transition from record k to record $k+1$ in a period. Finally, define δ_k to be the total share of players in record k that die in a period. The steady state equation for record k is

$$\delta_k + \tau_k = i_k, \tag{7}$$

which guarantees that the total inflow of players into record k precisely balances with the total outflow of players from record k.

Lemma 4. *Consider an equilibrium with total share of cooperators μ^C.*

1. $\delta_k = (1-\gamma)\mu_k$.
2. $i_0 = 1-\gamma$ *and* $i_{k+1} = \tau_k$ *for all* k.

Lemma 5. *Consider a steady state with total share of cooperators μ^C.*

1. *If k is a reciprocator record, then* $\tau_k = \gamma(1-(1-\varepsilon)\mu^C)\mu_k$.
2. *If k is an unconditional cooperator record, then* $\tau_k = \gamma\varepsilon\mu_k$.
3. *If k is a defector record, then* $\tau_k = \gamma\mu_k$.

Part 1 of Lemma 4 follows because each player dies in a given period with probability γ, independently of every other living player. Part 2 is a consequence of the count D's record-keeping system. Lemma 5 comes from Lemma 2 and a player's play and survival being independent.

It will also be useful to compute the probability that a player with record k increases her record to $k+1$ before dying. When k is a defector record, this probability simply equals γ, the probability of survival to the next period. When k is an unconditional cooperator record, this value is $\sum_{n=0}^{\infty}(\gamma(1-\varepsilon))^n\gamma\varepsilon$, since from Lemma2 the probability that a record k player becomes a record $k+1$ player exactly n periods from now is $(\gamma(1-\varepsilon))^n\gamma\varepsilon$. Straightforward algebra shows that

$$\sum_{n=0}^{\infty}(\gamma(1-\varepsilon))^n\gamma\varepsilon = \frac{\gamma\varepsilon}{1-\gamma(1-\varepsilon)} := \alpha(\gamma,\varepsilon).$$

Similarly, when k is a reciprocator record, the probability that a player with record k increases her record to $k+1$ before dying is $\beta(\gamma,\varepsilon,\mu^C)$, where $\beta : (0,1)\times(0,1)\times[0,1]\to(0,1)$ is given by

$$\beta(\gamma,\varepsilon,\mu^C) = \frac{\gamma(1-(1-\varepsilon)\mu^C)}{1-\gamma(1-\varepsilon)\mu^C}.$$[32]

These probabilities play a key role in deriving the continuation value function.

Lemma 6. *Consider a steady state with total share of cooperators μ^C. Consider the value function of a player with record k.*

1. *If k is a reciprocator record, $V_k = (1-\beta(\gamma,\varepsilon,\mu^C))\mu^C + \beta(\gamma,\varepsilon,\mu^C)V_{k+1}$.*
2. *If k is an unconditional cooperator record, $V_k = (1-\alpha(\gamma,\varepsilon))(\mu^C-\mu^D l) + \alpha(\gamma,\varepsilon)V_{k+1}$.*
3. *If k is a defector record, $V_k = (1-\gamma)\mu^{UC}(1+g) + \gamma V_{k+1}$.*

Proof. The expected flow payoff of a reciprocator is μ^C. Combining this with Lemma 2 yields $V_k = (1-\gamma)\mu^C + \gamma(1-\varepsilon)V_k + \gamma(1-(1-\varepsilon)\mu^C)V_{k+1}$. Solving for V_k gives Part 1 of the result.

The expected flow payoff of an unconditional cooperator is $\mu^C - \mu^D l$. Combining this with Lemma 2 yields $V_k = (1-\gamma)(\mu^C-\mu^D l) + \gamma(1-\varepsilon)V_k + \gamma\varepsilon V_{k+1}$. Solving for V_k gives Part 2 of the result.

The expected flow payoff of a defector is $\mu^{UC}(1+g)$. Combining this with Lemma 2 yields $V_k = (1-\gamma)\mu^{UC}(1+g) + \gamma V_{k+1}$. Solving for V_k gives Part 3 of the result. ■

[32]This comes from $\beta(\gamma,\varepsilon,\mu^C) = \sum_{n=0}^{\infty}(\gamma(1-\varepsilon)\mu^C)^n\gamma(1-(1-\varepsilon)\mu^C)$.

A.5 Proof of Theorem 3

The proof that any equilibrium that satisfies forgery-proofness and coordination-proofness corresponds to $GrimKL$ for some K, L proceeds by establishing the following two lemmas, the proofs of which are in OA.5.

Lemma 7. *In any equilibrium that satisfies forgery-proofness, there exists a record $\bar{k}$ such that a record k is an unconditional defector record iff $k \geq \bar{k}$.*

To see why this result holds, note that the $(C|C)_k$ constraint implies that $V_k - V_{k+1}$ is uniformly bounded from below by a strictly positive value for all records k at which a player is not an unconditional defector. Forgery-proofness then implies that there can be only finitely many records at which players are not unconditional defectors, because otherwise there would be some record k at which $V_k < 0$. Additionally, a defector record can never precede a cooperator record in a forgery-proof equilibrium. Otherwise, the defector would forge their record to skip to the next record at which some cooperation occurs.

Lemma 8. *In any equilibrium that satisfies forgery-proofness and coordination-proofness, there exists a record k^* such that all records $k < k^*$ are reciprocators and all records $k \geq k^*$ are either unconditional cooperators or unconditional defectors.*

Intuitively, since survival is uncertain, players prefer to obtain the higher flow payoff from reciprocation earlier than the lower flow payoff from unconditional cooperation, so if a UC record k was followed by an R record $k+1$, a player could gain by playing R at record k and UC at $k+1$.

These lemmmas show that only $GrimKL$ can satisfy both forgery-proofness and coordination-proofness. For the converse, note that Grimk KL satisfies Coordination Proofness. Moreover, if $GrimKL$ is an equilibrium, it must satisfy $V_k > V_{k+1}$ for all $k < K + L$, as otherwise some reciprocator or unconditional cooperator would prefer to defect. Since in addition V_k is constant for all $k \geq K + L$, it follows that V_k is everywhere non-increasing, so forgery-proofness is satisfied.

A.6 Proof of Theorem 4

Parts 1 and 2 of Theorem 4 come from the following proposition.

Proposition 4. *There is a GrimK equilibrium with share of cooperators* μ^C *if and only if the following conditions hold:*

1. *Feasibility:*

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

2. *Incentives:*

$$(C|C)_0 : \mu^C \in \left(\frac{1+g-\sqrt{(1+g)^2 - \frac{4}{1-\varepsilon}g}}{2}, \frac{1+g+\sqrt{(1+g)^2 - \frac{4}{1-\varepsilon}g}}{2} \right)$$
$$(D|D)_{K-1} : \mu^C < \frac{1}{\gamma(1-\varepsilon)} \frac{l}{1+l}.$$

The proof of Proposition 4 is in OA.6. The feasibility constraint comes from solving for the steady state record shares μ_k and using the fact that $\mu^C = \sum_{k=0}^{K-1} \mu_k$. The $(C|C)_0$ and $(D|D)_{K-1}$ incentive constraints come from computing the value functions for the *GrimK* strategy using Lemma 6 and then applying Lemma 3. It is without loss to restrict attention to the $(C|C)$ constraint for record 0 and the $(D|D)$ constraint for record $K-1$, as these two incentive constraints imply all of the others, because $V_k - V_{k+1}$ is increasing in k for $0 \leq k \leq K-1$. Intuitively, the incentive for reciprocators to play D when they should play C is greatest when $k = 0$ as then they are farthest away from the punishment phase. Likewise, the incentive for reciprocators to play C when they should play D is greatest when they are closest to the punishment phase, which is when their record is $k = K-1$.

To see how Part 2 of Theorem 4 comes from Proposition 4, note that

$$\lim_{\varepsilon \to 0} \frac{1+g-\sqrt{(1+g)^2 - \frac{4}{1-\varepsilon}g}}{2} = g$$

and

$$\lim_{\varepsilon\to 0}\lim_{\gamma\to 1}\frac{1}{\gamma(1-\varepsilon)}\frac{l}{1+l}=\frac{l}{1+l}.$$

Thus, it is impossible to satisfy both the $(C|C)_0$ and $(D|D)_{K-1}$ constraints when $g > l/(1+l)$, which is precisely when $l < g/(1-g)$. OA.7 gives a slightly modified argument that handles the case where $l = g/(1-g)$. To see how Part of 1 Theorem 4 comes from Proposition 4, note that, in addition to the above two limit results,

$$\lim_{\varepsilon\to 0}\frac{1+g+\sqrt{(1+g)^2-\frac{4}{1-\varepsilon}g}}{2}=1.$$

It follows that values of μ^C smaller than, but arbitrarily close to, $l/(1+l)$ satisfy the $(C|C)_0$ and $(D|D)_{K-1}$ constraints in the iterated limit. Thus, the only issue is the feasibility of μ^C as a steady-state level of cooperation. Because K must be an integer, some values of μ^C cannot be generated by any K, for given values of γ and ε. The following result shows that this "integer problem" become irrelevant in the iterated limit. Intuitively, any value of $\mu^C \in [0,1]$ can be approximated aritrarily closely by a feasible share of cooperators in a *GrimK* steady state as $\gamma \to 1$.

Lemma 9. *Fix $\varepsilon \in (0,1)$. For all $\Delta > 0$, there exists $\overline{\gamma} < 1$ such that, for all $\gamma > \overline{\gamma}$ and $\mu^C \in [0,1]$, there exists a $\hat{\mu}^C$ satisfying $|\hat{\mu}^C - \mu^C| < \Delta$ that satisfies the Feasibility constraint of Proposition 4 for some K.*

The proof of Lemma 9 is in OA.8.

A.7 Proof of Theorem 5

A.7.1 Proof Outline

Theorem 5 comes from the following characterization of *GrimKL* equilibria.

Proposition 5. *There is a GrimKL equilibrium with total share of cooperators μ^C, share of reciprocators μ^R, and share of unconditional cooperators μ^{UC} if and only if the* *following conditions hold:*

1. *Feasibility:*

$$\mu^C = 1 - \alpha(\gamma, \varepsilon)^L \beta(\gamma, \varepsilon, \mu^C)^K,$$
$$\mu^R = 1 - \beta(\gamma, \varepsilon, \mu^C)^K,$$
$$\mu^{UC} = (1 - \alpha(\gamma, \varepsilon)^L)\beta(\gamma, \varepsilon, \mu^C)^K.$$

2. *Incentives:*

$$(C|C)_0 : \frac{(1-\varepsilon)(1-\mu^C)}{1-(1-\varepsilon)\mu^C}\Big[\mu^R + \mu^{UC}(l-g)\Big] > g,$$
$$(D|D)_{K-1} : \frac{\gamma(1-\varepsilon)(1-\mu^C)}{1-\gamma(1-\varepsilon)\mu^C}\Big[\mu^R + \mu^{UC}(l-g)\Big] + \mu^R l < l,$$
$$(C|D)_K : \frac{(1-\varepsilon)(1-\mu^C)}{1-(1-\varepsilon)\mu^C}\Big[\mu^R + \mu^{UC}(l-g)\Big] + \mu^R l > l.$$

The proof, which is similar to that of Proposition 4, is in OA.9. In particular, the feasibility constraints come from calculating the steady state shares μ_k for the $GrimKL$ strategy and then setting $\mu^C = \sum_{k=0}^{K+L-1} \mu_k$, $\mu^R = \sum_{k=0}^{K-1} \mu_k$, and $\mu^{UC} = \sum_{k=K}^{K+L-1} \mu_k$. Moreover, the $(C|C)_0$, $(D|D)_{K-1}$, and $(C|D)_K$ incentive constraints come from computing the value functions for the $GrimKL$ strategy using Lemma 6 and then applying Lemma 3.

Part 3 of Theorem 5 follows from Proposition 5 because $(C|C)_0$ and $(C|D)_K$ are impossible to satisfy simultaneously when $g \geq 1$ and $l \leq g(g+1)$, regardless of γ and ε. Part 2 also follows from Proposition 5, because, as γ approaches 1, the left-hand side of the $(D|D)_{K-1}$ constraint becomes identical to the left-hand side of the $(C|D)_K$ constraint, so both constraints must hold with equality in this limit. Combining this with $(C|C)_0$ shows that $\mu^R \leq 1 - g/l$ must be satisfied for all $GrimKL$ equilibria with $L > 0$ in the $\gamma \to 1$ limit. However, neither the $(C|C)_0$ nor the $(D|D)_{K-1}$ constraints can be satisfied with $\mu^R \leq 1 - g/l$ when $l \leq g(g+1)$. The formal proofs of Parts 2 and 3 of Theorem 5 are given in A.7.2.

Proving that Part 1 of Theorem 5 follows from Proposition 5 is more involved than proving that Part 1 of Theorem 4 follows from Proposition 4. The demonstration proceeds by first identifying a target level of cooperation for fixed ε. The greatest difficulty in the proof involves showing that there are feasible profiles satisfying the incentive constraints where the level of cooperation actually attains this target as γ approaches 1. The proof then shows that that this level of cooperation approaches 1 as ε approaches 0.

A.7.2 Proof of Parts 2 and 3 of Theorem 5

Proof of Part 3 of Theorem 5. Fix $\varepsilon \in (0,1)$, $g \geq 1$, and $l \leq g(g+1)$. Since

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

for all $\mu^C \in [0,1]$, the $(C|C)_0$ constraint requires

$$\mu^R + \mu^{UC}(l-g) > g, \tag{8}$$

and the $(C|D)_K$ constraint requires

$$\mu^R(1+l) + \mu^{UC}(l-g) > l. \tag{9}$$

Since $\mu^R \geq 0$, $\mu^{UC} \geq 0$, and $\mu^R + \mu^{UC} \leq 1$, Inequality 8 requires $l > 2g$. Note that $l > 2g$ and $l \leq g(g+1)$ cannot be jointly satisfied when $g = 1$, and hence the $(C|C)_0$ constraint cannot be satisfied when $g = 1$. As equilibria with unconditional cooperators but not reciprocators are impossible, it follows that, when $g = 1$, there is no equilibrium with cooperators, so $\overline{\mu}^C_{KL}(\gamma, \varepsilon) = 0$ when $g = 1$. Now, consider $g > 1$. Inequality 8 additionally requires $\mu^R < \frac{l-2g}{l-g-1}$.

Because $(l-2g)/(l-g-1)$ is strictly increasing in l for $l > 2g$, it follows that

$$\mu^R < \frac{g(g+1)-2g}{g(g+1)-g-1} = \frac{g}{g+1}, \tag{10}$$

since $l \leq g(g+1)$ by assumption. Similarly, Inequality 9 requires

$$\mu^R > \frac{g}{g+1}. \tag{11}$$

However, Inequalities 10 and 11 are mutually incompatible, so the $(C|C)_0$ and $(C|D)_K$ constraints cannot be jointly satisfied. Therefore, there is no $GrimKL$ equilibrium with both reciprocators and unconditional cooperators. As Inequality 8 cannot be satisfied with $\mu^{UC} = 0$ and equilibria with unconditional cooperators but not reciprocators are impossible, it follows that there is no equilibrium with cooperators. Consequently, $\overline{\mu}^C_{KL}(\gamma, \varepsilon) = 0$ when $g > 1$. ■

Proof of Part 2 of Theorem 5. We show that $\limsup_{\gamma \to 1} \overline{\mu}^C_{KL}(\gamma, \varepsilon) = 0$ for all $\varepsilon \in (0,1)$ when $l \leq g(g+1)$. Suppose otherwise that $\limsup_{\gamma \to 1} \overline{\mu}^C_{KL}(\gamma, \varepsilon) > 0$. Then there is some $\gamma_n \to 1$ and a sequence of associated equilibria with shares (μ^R_n, μ^{UC}_n), such that (μ^R_n, μ^{UC}_n) converges to (μ^R, μ^{UC}) with $\mu^C = \mu^R + \mu^{UC} > 0$. Theorem 4 implies that $\limsup_{\gamma \to 1} \overline{\mu}^C_K(\gamma, \varepsilon) = 0$ for $l \leq g(g+1)$, so such a sequence must satisfy the $(C|C)_0$, $(C|D)_{K-1}$, and $(D|D)_K$ constraints for each corresponding γ_n. Taking the limit of these constraints as $n \to \infty$ shows that (μ^R, μ^{UC}) must satisfy the following "limit" constraints.

$$\begin{aligned}
\text{Limit } (C|C)_0 : \ & \frac{(1-\varepsilon)(1-\mu^C)}{1-(1-\varepsilon)\mu^C}(\mu^R + \mu^{UC}(l-g)) \geq g, \\
\text{Limit } (D|D)_{K-1} : \ & \frac{(1-\varepsilon)(1-\mu^C)}{1-(1-\varepsilon)\mu^C}(\mu^R + \mu^{UC}(l-g)) + \mu^R l \leq l, \\
\text{Limit } (C|D)_K : \ & \frac{(1-\varepsilon)(1-\mu^C)}{1-(1-\varepsilon)\mu^C}(\mu^R + \mu^{UC}(l-g)) + \mu^R l \geq l.
\end{aligned} \tag{12}$$

The limit $(D|D)_{K-1}$ and $(C|D)_K$ constraints together imply that

$$\frac{(1-\varepsilon)(1-\mu^C)}{1-(1-\varepsilon)\mu^C}(\mu^R+\mu^{UC}(l-g))+\mu^R l = l.$$

Along with the limit $(C|C)_0$ constraint, this implies that $\mu^R \leq 1-g/l$.

For any $\varepsilon \in (0,1)$,

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

for all $\mu^C \in [0,1]$. Therefore, the Limit $(C|D)_K$ constraint requires

$$\mu^R(1+l)+\mu^{UC}(l-g) > l. \tag{13}$$

Since $0 \leq \mu^R \leq 1-g/l$, $\mu^{UC} \geq 0$, and $\mu^R+\mu^{UC} \leq 1$, it follows that

$$\mu^R(1+l)+\mu^{UC}(l-g) \leq (1-\frac{g}{l})(1+l)+\frac{g}{l}(l-g) = l+\frac{1}{l}(l-g(g+1)).$$

Since $l \leq g(g+1)$, this implies that $\mu^R(1+l)+\mu^C(l-g) \leq l$, a contradiction. ■

A.7.3 Proof of Part 1 of Theorem 5

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

$$\frac{(1-\varepsilon)(1-\mu^C)}{1-(1-\varepsilon)\mu^C}\Big[(l-g)\mu^C+(1+g-l)\mu^R\Big]+l\mu^R = l \tag{14}$$

and the function $h(\varepsilon,\mu^R)$ defined by $h(\varepsilon,\mu^R) = \max\{\mu^C \in [0,1] : \mu^C \text{ solves Equation 14}\}$. If $h(\varepsilon,\mu^R)$ is well-defined, it gives the maximum level of cooperation for the given μ^R and ε that satisfies the $\gamma \to 1$ "limit" constraints of Equation 12. Straightforward calculations show that, for any $\mu^R \in (g/(1+g), 1-g/l]$, $h(\varepsilon,\mu^R)$ is well-defined for sufficiently small and positive ε, and that

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

An immediate implication of this is $\lim_{\varepsilon\to 0} h(\varepsilon, \mu^R) = 1$. Combining this with the following two lemmas proves Part 1 of Theorem 5.

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

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

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

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

Lemma 10. *Fix $\mu^R \in (g/(1+g), 1-g/l]$. If $|1+\kappa(\mu^R)| > \iota(\mu^R)$, then there exists some $\overline{\varepsilon} > 0$, such that $\liminf_{\gamma\to 1} \overline{\mu}^C_{KL}(\gamma, \varepsilon) \geq h(\varepsilon, \mu^R)$ for $\varepsilon < \overline{\varepsilon}$.*

Lemma 11. *Suppose that $l > g(g+1)$. Some $\mu^R \in (g/(1+g), 1-g/l]$ satisfies $|1+\kappa(\mu^R)| > \iota(\mu^R)$ if $l > \max\{g(g+1), b(g)\}$.*

OA.10.1 presents the proof of Lemma 10. It makes heavy use of the inverse function theorem and other tools of differential calculus to show that, when $|1+\kappa(\mu^R)| > \iota(\mu^R)$, for sufficiently small ε, any neighborhood of $(h(\varepsilon, \mu^R), \mu^R)$ can be approached by feasible profiles for sufficiently high γ. The proof of Lemma 11 is in OA.10.2.

Online Appendix for "Robust Cooperation with First-Order Information"

Daniel Clark, Drew Fudenberg, and Alexander Wolitzky

April 19, 2019

OA.1 Proofs of Results for Theorem 1

Recall that $R(t)$ is the set of feasible records for a player of age t, and
$\bar{M} = \left\{ \mu \in M : \sum_{k \in R(t)} \mu_k \leq \gamma^t \ \forall t \in \mathbb{N} \right\}$.

Claim 1. *$\bar{M}$ is compact 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 it is sequentially compact. Consider a sequence $\{\mu^n\}_{n\in\mathbb{N}}$ of $\mu^n \in \bar{M}$. By a standard diagonalization argument, there exists some $\tilde{\mu} \in [0,1]^{\mathbb{Z}}$ and some subsequence $\{\mu^{n_m}\}_{m\in\mathbb{N}}$ such that $\lim_{m\to\infty} \mu_k^{n_m} = \tilde{\mu}_k$ for all $k \in \mathbb{Z}$. This, along with the fact that

$$\sum_{t'<t} \sum_{k \in R(t')} \mu_k^{n_m} \geq 1 - \frac{\gamma^t}{1-\gamma}$$

*We thank V Bhaskar, Glenn Ellison, 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 finanical support.

for all $t \in \mathbb{N}$ and $m \in \mathbb{N}$ implies that

$$\sum_{t'<t} \sum_{k \in R(t')} \tilde{\mu}_k \geq 1 - \frac{\gamma^t}{1-\gamma}$$

for all $t \in \mathbb{N}$. Thus, $\sum_{k \in \mathbb{Z}} \tilde{\mu}_k = 1$ since $\gamma^t \to 0$ as $t \to \infty$, so $\tilde{\mu} \in M$. Additionally, since

$$\sum_{k \in R(t)} \mu_k^{n_m} \leq \gamma^t$$

for all $t \in \mathbb{N}$ and $m \in \mathbb{N}$, it follows that

$$\sum_{k \in R(t)} \tilde{\mu}_k \leq \gamma^t$$

for all $t \in \mathbb{N}$, so $\tilde{\mu} \in \bar{M}$.

Now, we show that $\lim_{m \to \infty} \mu^{n_m} = \tilde{\mu}$. Fix $\varepsilon > 0$. Let $T \in \mathbb{N}$ be such that $\gamma^T < \varepsilon$, and let $M \in \mathbb{N}$ be such that $|\mu_k^{n_m} - \tilde{\mu}_k| < \varepsilon$ for all $k \in R(t)$, $t < T$ and $m > M$. Thus,

$$\sup_{k \in \mathbb{Z}} |\mu_k^{n_m} - \tilde{\mu}_k| < \varepsilon$$

for all $m > M$. ■

Claim 2. *$f_{r,s}$ maps $\bar{M}$ to itself.*

Proof of Claim 2. By the definition of $R(t)$, for all $\mu \in M$ and all $t \in \mathbb{N}$, we have $\gamma \sum_{k \in R(t)} \mu_k = \sum_{k \in R(t+1)} f_{r,s}(\mu)[k]$. Hence, if $\sum_{k \in R(t)} \mu_k \leq \gamma^t$ for all $t \in \mathbb{N}$ then $\sum_{k \in R(t+1)} f_{r,s}(\mu)[k] \leq \gamma^{t+1}$ for all $t \in \mathbb{N}$ (and, trivially, $\sum_{k \in R(0)} f_{r,s}(\mu)[k] \leq 1$). That is, $f_{r,s}$ maps $\bar{M}$ to itself. ■

Claim 3. *$f_{r,s}$ is continuous in the* sup *norm topology.*

Proof of Claim 3. Given a vector $\mu \in \bar{M}$, define the vector $\mu_{|T} \in \bar{M}_T$ by setting $\mu_{t|T} = \mu_\tau$ for $t \in \bigcup_{t \leq T} R(t)$ and $\mu_{t|T} = 0$ for $t \notin \bigcup_{t \leq T} R(t)$. For each T, define the function $f_{r,s|T} : \bar{M} \to \bar{M}$ by setting $f_{r,s|T}(\mu) = f_{r,s}(\mu)_{|T}$ for each $\mu \in \bar{M}$. Note that $f_{r,s|T}(\mu) =$

$f_{r,s|T}(\mu')$ whenever $\mu_{|\bigcup_{t\le T} R(t)} = \mu'_{|\bigcup_{t\le T} R(t)}$. Thus, $f_{r,s|T}$ can equivalently be viewed as a function from $\bar{M}_T$ to $\bar{M}$, and as a polynomial function it is continuous on $\bar{M}_T$, and hence on $\bar{M}$. Moreover, since the mass on records outside $\bigcup_{t\le T} R(t)$ goes to 0 as $T\to\infty$, for any $\varepsilon>0$ there exists T such that $|f_{r,s}(\mu) - f_{r,s}(\mu)_{|T}| < 2\varepsilon$ for all $\mu \in \bar{M}$. Hence, $f_{r,s}$ is also continuous on $\bar{M}$. ■

OA.2 Proof of Proposition 1

Proposition 1. *Limit efficiency is attainable in strict equilibrium whenever the prisoner's dilemma is mild ($g<1$).*

Assume $g<1$. Fix any rational number ρ satisfying $g<\rho<\min\{l,1\}$. Let m and n be integers such that $m\ge n>0$ and $n/m=\rho$.

We consider a strategy with $m+n$ *phases*, where a player is in phase j whenever her record equals $j-1 \bmod m+n$. The first m phases, denoted G_1 through G_m, are *good phases*, and the remaining n phases, denoted B_{m+1} through B_{m+n}, are *bad phases*. A player is a reciprocator while in a good phase and a defector while in a bad phase. We denote the share of players in phase G_j by μ_{G_j} and the share of players in phase B_j by μ_{B_j}. Consequently, the total share of cooperators is $\mu^C=\sum_{j=1}^m \mu_{G_j}$ and the total share of defectors is $\mu^D=\sum_{j=m+1}^{m+n}\mu_{B_j} = 1-\mu^C$.

We first prove that under this strategy the share of cooperators μ^C converges to 1 in the iterated limit where γ approaches 1 and then ε approaches 0.[1] We then prove that the strategy does in fact give strict equilibria. This result holds for all $\rho<1$, so for any g and l such that $g<\min\{l,1\}$, there is a strategy that obtains limit efficiency. This proves Proposition 1.

Lemma 12. $\lim_{\varepsilon\to0}\lim_{\gamma\to1}\mu^C=1$.

Before proving the lemma, we give a heuristic argument. As $\gamma\to1$, the mass μ^C of reciprocators will be approximately equally distributed among the first m phases, while

[1]Here and subsequently we do not track the dependence of endogenous objects like μ^C on γ and ε in the notation.

the mass $1-\mu^C$ of defectors will be approximately equally distributed among the next n phases. Following $\gamma \to 1$, as $\varepsilon \to 0$, the flow from phase $m+n$ to phase 1 (the "inflow into cooperation") is approximately $(1-\mu^C)/n$, while the flow from phase m to phase $m+1$ (the "inflow into defection") is approximately $(1-\mu^C)(\mu^C/m)$. If these flows were equal for some $\mu^C < 1$, this would imply $\mu^C = m/n < 1$. But this contradicts the fact that $m \geq n$ as there are more good phases than bad phases. Therefore, a steady state requires that $\mu^C = 1$ in the iterated limit.

Proof of Lemma 12. Since the first m phases all correspond to reciprocator records, Lemmas 4 and 5 imply that $i_{G_j} = \gamma(1-(1-\varepsilon)\mu^C)\mu_{G_{j-1}}$ and $\tau_{G_j} = \gamma(1-(1-\varepsilon)\mu^C)\mu_{G_j}$ for all $1 < j \leq m$. By Equation 7, $\mu_{G_j} = \beta(\gamma,\varepsilon,\mu^C)\mu_{G_{j-1}}$, so induction gives

$$\mu_{G_j} = \beta(\gamma,\varepsilon,\mu^C)^{j-1}\mu_{G_1} \tag{15}$$

for $1 \leq i \leq m$.

Since the phase m corresponds to reciprocator records and phase $m+1$ corresponds to defector records, Lemmas 4 and 5 give $i_{B_{m+1}} = \gamma(1-(1-\varepsilon)\mu^C)\mu^{G_m}$ and $\tau_{B_{m+1}} = \gamma\mu_{B_{m+1}}$, so Equation 7 implies

$$\begin{aligned} \mu_{B_{m+1}} &= \gamma(1-(1-\varepsilon)\mu^C)\mu^{G_m} \\ &= \beta(\gamma,\varepsilon,\mu^C)^m(1-\gamma(1-\varepsilon)\mu^C)\mu_{G_1}. \end{aligned} \tag{16}$$

Since the last n phases all correspond to defector records, Lemmas 4 and 5 give $i_{B_j} = \gamma\mu_{B_{j-1}}$ and $\tau_{B_j} = \gamma\mu_{B_j}$ for $m < j \leq m+n$. Thus, Equation 7 implies that $\mu_{B_j} = \gamma\mu_{B_{j-1}}$, so induction, combined with Equation 16, gives

$$\begin{aligned} \mu_{B_{m+n}} &= \gamma^{n-1}\mu_{B_{m+1}} \\ &= \gamma^{n-1}\beta(\gamma,\varepsilon,\mu^C)^m(1-\gamma(1-\varepsilon)\mu^C)\mu_{G_1}. \end{aligned} \tag{17}$$

Finally, since phase 1 corresponds to reciprocator records and phase $m+n$ corresponds to defector records, Lemmas 4 and 5 give $i_{G_1} = 1-\gamma+\gamma\mu^{B_{m+n}}$ and $\tau_{G_1} =$

$\gamma(1-(1-\varepsilon)\mu^C)\mu_{G_1}$, so Equations 7 and 17 imply

$$\begin{aligned}\mu^{G1} &= \frac{1-\gamma}{1-\gamma(1-\varepsilon)\mu^C} + \frac{\gamma}{1-\gamma(1-\varepsilon)\mu^C}\mu_{B_{m+n}} \\ &= 1-\beta(\gamma,\varepsilon,\mu^C) + \gamma^n\beta(\gamma,\varepsilon,\mu^C)^m\mu^{G1}.\end{aligned}$$

Solving this for μ_{G_1} gives

$$\mu_{G_1} = \frac{1-\beta(\gamma,\varepsilon,\mu^C)}{1-\gamma^n\beta(\gamma,\varepsilon,\mu^C)^m}. \tag{18}$$

Equations 15 and 18 together imply that

$$\begin{aligned}\sum_{j=1}^{m}\mu_{G_j} &= \sum_{j=1}^{m}\beta(\gamma,\varepsilon,\mu^C)^{j-1}\Big(\frac{1-\beta(\gamma,\varepsilon,\mu^C)}{1-\gamma^n\beta(\gamma,\varepsilon,\mu^C)^m}\Big) \\ &= \frac{1-\beta(\gamma,\varepsilon,\mu^C)^m}{1-\gamma^n\beta(\gamma,\varepsilon,\mu^C)^m}.\end{aligned}$$

Therefore,

$$\mu^C = \frac{1-\beta(\gamma,\varepsilon,\mu^C)^m}{1-\gamma^n\beta(\gamma,\varepsilon,\mu^C)^m}. \tag{19}$$

Consider the function $f:[0,1]\times(0,1)\times[0,1]\to\mathbb{R}$ given by

$$f(\gamma,\varepsilon,\mu^C) = \begin{cases}\frac{1-\beta(\gamma,\varepsilon,\mu^C)^m}{1-\gamma^n\beta(\gamma,\varepsilon,\mu^C)^m} & \text{if } \gamma<1 \\ \frac{1}{1+\rho(1-(1-\varepsilon)\mu^C)} & \text{if } \gamma=1\end{cases}.$$

This function extends $(1-\beta(\gamma,\varepsilon,\mu^C)^m)/(1-\gamma^n\beta(\gamma,\varepsilon,\mu^C)^m)$ to $\gamma=1$ and, for any fixed value of $\varepsilon\in(0,1)$, is a continuous function of $(\gamma,\mu^C)\in[0,1]\times[0,1]$, which can be shown using L'Hôpital's rule.

Therefore, for fixed ε, any limit point of any sequence of steady state μ^C as $\gamma\to1$ must satisfy

$$\begin{aligned}\mu^C &= f(1,\varepsilon,\mu^C) \\ &= \frac{1}{1+\rho(1-(1-\varepsilon)\mu^C)}.\end{aligned}$$

The only such $\mu^C \in [0,1]$ that satisfies this equation for ε is

$$\overline{\mu}^C(\varepsilon) = \frac{1+\rho-\sqrt{(1+\rho)^2-4(1-\varepsilon)\rho}}{2(1-\varepsilon)\rho},$$

so $\lim_{\gamma\to 1}\mu^C = \overline{\mu}^C(\varepsilon)$ for all ε. Lemma 12 follows since $\lim_{\varepsilon\to 0}\overline{\mu}^C(\varepsilon) = 1$. ■

The next lemma formalizes the idea that, on average, adding a D to one's record while in a good phase leads to an extra ρ periods of punishment. Let V_{G_j} denote the value function of a player in phase G_j and V_{B_j} denote the value function of a player in phase B_j.

Lemma 13. *The following iterated limits hold:*

$$\lim_{\varepsilon\to 0}\lim_{\gamma\to 1}\frac{V_{G_j}-V_{G_j}}{1-\gamma} = \rho \text{ for } 1 \le j < m \text{ and } \lim_{\varepsilon\to 0}\lim_{\gamma\to 1}\frac{V_{G_m}-V_{B_{m+1}}}{1-\gamma} = \rho.$$

By Lemma 3, this implies that the $(C|C)$ and $(D|D)$ constraints for a player in a good phase are satisfied in the iterated limit when $g < \rho < l$. The incentives for players in bad phases are trivial, because the value function in the next phase is larger than the value function in the current phase whenever the current phase is bad. This implies that playing D while in a bad phase maximizes both the flow payoff and the continuation payoff. Therefore, Lemma 3, along with Lemma 12, suffices to prove Proposition 1.

Proof of Lemma 13. By Lemma 6,

$$\frac{V_{G_j}-V_{G_{j+1}}}{1-\gamma} = \frac{\mu^C - V_{G_{j+1}}}{1-\gamma(1-\varepsilon)\mu^C}$$

for all $1 \le j < m$ and

$$\frac{V_{G_m}-V_{B_{m+1}}}{1-\gamma} = \frac{\mu^C - V_{B_{m+1}}}{1-\gamma(1-\varepsilon)\mu^C}.$$

Lemma 6 also implies that the value functions in all phases converge to the same value as $\gamma \to 1$ for fixed ε. Therefore, it suffices to show that

$$\lim_{\varepsilon\to 0}\lim_{\gamma\to 1}\frac{\mu^C - V_{G_1}}{1-\gamma(1-\varepsilon)\mu^C} = \rho.$$

Since the first m phases all correspond to reciprocator records, combining Lemma 6 with induction gives

$$V_{G_1} = (1-\beta(\gamma,\varepsilon,\mu^C)^m)\mu^C + \beta(\gamma,\varepsilon,\mu^C)^m V_{B_{m+1}}.$$

Likewise, since the last n phases all correpsond to defector records, Lemma 6, along with the fact that $\mu^{UC} = 0$ in the strategies considered, implies

$$V_{B_{m+1}} = \gamma^n V_{G_1}.$$

Combining these equations gives

$$V_{G_1} = (1-\beta(\gamma,\varepsilon,\mu^C)^m)\mu^C + \gamma^n\beta(\gamma,\varepsilon,\mu^C)V_{G_1},$$

and solving this for V_{G_1} renders

$$V_{G_1} = \Big(\frac{1-\beta(\gamma,\varepsilon,\mu^C)^m}{1-\gamma^n\beta(\gamma,\varepsilon,\mu^C)^m}\Big)\mu^C.$$

Therefore,

$$\frac{\mu^C - V_{G_1}}{1-\gamma(1-\varepsilon)\mu^C} = \Big(\frac{(1-\gamma^n)\beta(\gamma,\varepsilon,\mu^C)^m}{(1-\gamma(1-\varepsilon)\mu^C)(1-\gamma^n\beta(\gamma,\varepsilon,\mu^C)^m)}\Big)\mu^C.$$

Since $\lim_{\varepsilon\to 0}\lim_{\gamma\to 1}\mu^C = 1$ by Lemma 12, we need only show that

$$\lim_{\varepsilon\to 0}\lim_{\gamma\to 1}\frac{(1-\gamma^n)\beta(\gamma,\varepsilon,\mu^C)^m}{(1-\gamma(1-\varepsilon)\mu^C)(1-\gamma^n\beta(\gamma,\varepsilon,\mu^C)^m)} = \rho. \tag{20}$$

Consider the function $\tilde{f} : [0,1] \times (0,1) \times [0,1] \to \mathbb{R}$ given by

$$\tilde{f}(\gamma, \varepsilon, \mu^C) = \begin{cases} \frac{(1-\gamma^n)\beta(\gamma,\varepsilon,\mu^C)^m}{(1-\gamma(1-\varepsilon)\mu^C)(1-\gamma^n\beta(\gamma,\varepsilon,\mu^C)^m)} & \text{if } \gamma < 1 \\ \frac{\rho}{1+\rho(1-(1-\varepsilon)\mu^C)} & \text{if } \gamma = 1 \end{cases}.$$

For any fixed value of ε, this function is a continuous function of $(\gamma, \mu^C) \in [0,1] \times [0,1]$, which can be shown using L'Hôpital's rule. Therefore,

$$\begin{aligned} \lim_{\varepsilon \to 0} \lim_{\gamma \to 1} \frac{(1-\gamma^n)\beta(\gamma,\varepsilon,\mu^C)^m}{(1-\gamma(1-\varepsilon)\mu^C)(1-\gamma^n\beta(\gamma,\varepsilon,\mu^C)^m)} &= \lim_{\varepsilon \to 0} \lim_{\gamma \to 1} \tilde{f}(\gamma, \varepsilon, \mu^C) \\ &= \lim_{\varepsilon \to 0} \frac{\rho}{1+\rho(1-(1-\varepsilon)\overline{\mu}^C(\varepsilon))} \\ &= \rho, \end{aligned}$$

where the second and third equalities follow from the facts established in the proof of Lemma 12 that $\lim_{\gamma \to 1} \mu^C = \overline{\mu}^C(\varepsilon)$ and $\lim_{\varepsilon \to 0} \overline{\mu}^C(\varepsilon) = 1$. Thus, the limit in Equation 20 is satisfied, and Lemma 13 follows. ■

OA.3 Extension of Proposition 3

Proposition 6. *When records count C's, the unique strict equilibrium remains* Always Defect *even if plays of D are mis-recorded as C with probability $\varepsilon' > 0$.*

Proof. It suffices to show that $V_k \leq V_{k+1}$ for all $k \in \mathbb{N}$ and apply the proof of Proposition 3. So suppose $V_k > V_{k+1}$ for some k. Then a player with record k always plays D: this maximizes both her flow payoff and her continuation payoff. Since there are no anti-reciprocators, this implies that only unconditional cooperators play C against a player with record k.

Now, let π_k denote the flow payoff of a player with record k, i.e. a player who always plays D while facing C from unconditional cooperators and D from everyone else. We claim that $\pi_k \leq V_{k'}$ for every record k', and in particular $\pi_k \leq V_{k+1}$. To see this, note that a player who always plays D gets at least π_k in every period (as unconditional cooperators always play C against her, and depending on her record maybe other players do, too). So if $\pi_k > V_{k'}$ then a player with record k' would do strictly better to play D forever.

We now have $V_k > V_{k+1}$ and $\pi_k \leq V_{k+1}$. But V_k is a weighted average of π_k and V_{k+1}, a contradiction. ■

OA.4 Erasure-Proofness

Proposition 7. *For any record-keeping system, the only erasure-proof equilibrium is* Always Defect*.*

Proof. Let $\bar{\pi}$ be the supremum of the flow payoffs earned at any record. We claim first that if an equilibrium prescribes cooperation at any record (whether or not it is erasure-proof), then no player with an unconditional defector record can earn flow payoffs within $\frac{1-\gamma}{\gamma}g$ of $\bar{\pi}$. This follows from the fact that there are no anti-reciprocators, so at any record a player who deviates to D will receive flow payoff no less than that of players who are prescribed to play D. Hence, at any record where cooperation is prescribed, deviating to always playing D increases a player's instantaneous payoff by at least $(1-\gamma)g$ and reduces her continuation payoff by at most $\frac{1-\gamma}{\gamma}g$, and thus constitutes a profitable deviation.

Now let k be a record where the flow payoff is within $\frac{1-\gamma}{\gamma}g$ of $\bar{\pi}$. Then any erasure-proof equilibrium must prescribe unconditional defection at record k, as playing D and erasing the record update to keep one's record fixed at k increases one's instantaneous payoff by at least $(1-\gamma)g$ and reduces one's continuation payoff by at most $\frac{1-\gamma}{\gamma}g$.

Combining these two observations implies that players must defect at every record in any erasure-proof equilibrium. ■

OA.5 Proofs of Results for Theorem 3

Lemma 14. *In every non-trivial equilibrium, $\pi_R > \pi_D$. In every equilibrium with unconditional cooperators, $\pi_R > \pi_{UC} > \pi_D$.*

Proof of Lemma 14. $\pi_R > \pi_{UC}$ follows from $\mu^D > 0$. To see that $\mu^D > 0$, note that, if $\mu^D = 0$, then every player would face only cooperators for the duration of her lifetime regardless of her history of play. However, then every player would defect in every period, a contradiction.

Next, if $\pi_D \geq \pi_{UC}$, then unconditional cooperators would receive the lowest flow payoff of any class. Since V_k is a convex combination of $\pi_{k'}$ for $k' \geq k$ and a player's record remains constant when she plays C, this implies that a player at any unconditional cooperator record would do strictly better by playing D until her record changes, a contradiction.

Hence, $\pi_R > \pi_{UC} > \pi_D$ in any equilibrium with unconditional cooperators. A similar argument implies $\pi_R > \pi_D$ in any non-trivial equilibrium without unconditional cooperators. ■

Lemma 7. *In any equilibrium that satisfies forgery-proofness, there exists a record $\bar{k}$ such that a record k is an unconditional defector record iff $k \geq \bar{k}$.*

Proof of Lemma 7. First, we establish that there must be some cutoff record after which a player is always an unconditional defector. Note that for any record k at which a player is not an unconditional defector, the $(C|C)_k$ constraint requires that

$$V_k - V_{k+1} > \frac{1-\gamma}{\gamma(1-\varepsilon)}g.$$

Thus, if in a forgery-proof equilibrium where the value function is everywhere non-increasing, there were infinitely many records at which a player is not an unconditional defector, there would be some k for which $V_k < 0$, which is impossible in equilibrium.

We now establish that unconditional defector records can only be followed by other unconditional defector records. The reason for this is that otherwise, there would be some unconditional defector record k at which a player would strictly prefer to inflate her record until she reaches the next record at which some cooperation occurs, which would violate forgery proofness. ■

Lemma 8. *In any equilibrium that satisfies forgery-proofness and coordination-proofness, there exists a record k^* such that all records $k < k^*$ are reciprocators and all records $k \geq k^*$ are either unconditional cooperators or unconditional defectors.*

Proof of Lemma 8. From Lemma 7, there exists some record $\overline{k}$ such that all records $k \geq \overline{k}$ are unconditional defector records and all records $k < \overline{k}$ are either reciprocator or unconditional cooperator records. Suppose that there are m records at which a player is a reciprocator and $\overline{k} - m$ records at which a player is an unconditional cooperator. We must show that the first m records, $0 \leq k \leq m-1$, correspond to R while the next $\overline{k} - m$ records, $m \leq k \leq \overline{k} - 1$, correspond to UC. Note that this is vacuously true if $m = 0$ or $m = \overline{k}$. We now show that it is true for $0 < m < \overline{k}$. Suppose towards a contradiction that there exists some k satisfying $0 \leq k < \overline{k} - 1$ which corresponds to UC and is such that $k+1$ corresponds to R. By the familiar recursive expressions for a player's expected payoff as a function of their record, a record $k+1$ player's expected payoff given by the strategy profile is $V_{k+1} = (1 - \beta(\gamma, \varepsilon, \mu^C))\pi_R + \beta(\gamma, \varepsilon, \mu^C)V_{k+2}$, and a record k player's expected payoff given by the strategy profile is

$$\begin{aligned} V_k &= (1 - \alpha(\gamma, \varepsilon))\pi_{UC} + \alpha(\gamma, \varepsilon)V_{k+1} \\ &= \alpha(\gamma, \varepsilon)(1 - \beta(\gamma, \varepsilon, \mu^C))\pi_R + (1 - \alpha(\gamma, \varepsilon))\pi_{UC} + \alpha(\gamma, \varepsilon)\beta(\gamma, \varepsilon, \mu^C)V_{k+2}, \end{aligned}$$

where V_{k+2} is the player's expected payoff upon reaching record $k+2$. However, suppose instead that the player changed their strategy so that she plays according to R at record k and according to UC at record $k+1$, but otherwise keeps her strategy the same. Then the player's expected payoff upon reaching record $k+1$, which we denote by $\tilde{V}_{k+1}$, would be $\tilde{V}_{k+1} = (1 - \alpha(\gamma, \varepsilon))\pi_{UC} + \alpha(\gamma, \varepsilon)V_{k+2}$, and the player's expected payoff upon reaching record k, which we denote by $\tilde{V}_k$, would be

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

Note that $\tilde{V}_k - V_k = (1-\alpha(\gamma,\varepsilon))(1-\beta(\gamma,\varepsilon,\mu^C))(\pi_R - \pi_{UC}) > 0$ where the inequality follows because $\pi_R > \pi_{UC}$ by Lemma 14. Thus the profile is not an equilibrium. ■

OA.6 Proof of Proposition 4

Proposition 4. *There is a GrimK equilibrium with share of cooperators μ^C if and only if the following conditions hold:*

1. *Feasibility:*
$$\mu^C = 1 - \beta(\gamma,\varepsilon,\mu^C)^K.$$

2. *Incentives:*
$$\begin{aligned}(C|C)_0 &: \mu^C \in \left(\frac{1+g-\sqrt{(1+g)^2 - \frac{4}{1-\varepsilon}g}}{2}, \frac{1+g+\sqrt{(1+g)^2 - \frac{4}{1-\varepsilon}g}}{2}\right), \\ (D|D)_{K-1} &: \mu^C < \frac{1}{\gamma(1-\varepsilon)}\frac{l}{1+l}.\end{aligned}$$

Lemma 15. *In a GrimK steady state with total share of cooperators μ^C,*

$$\mu_k = \begin{cases} \beta(\gamma,\varepsilon,\mu^C)^k(1-\beta(\gamma,\varepsilon,\mu^C)) & \text{if } 0 \le k \le K-1 \\ \gamma^{k-K}\beta(\gamma,\varepsilon,\mu^C)^K(1-\gamma) & \text{if } k \ge K \end{cases}. \tag{21}$$

Moreover, μ^C satisfies the equation

$$\mu^C = 1 - \beta(\gamma,\varepsilon,\mu^C)^K. \tag{22}$$

Proof. Since $i_0 = 1-\gamma$ and $\tau_0 = \gamma(1-(1-\varepsilon)\mu^C)\mu_0$, Equation 7 implies $\mu_0 = \frac{1-\gamma}{1-\gamma(1-\varepsilon)\mu^C} = 1-\beta(\gamma,\varepsilon,\mu^C)$. Moreover, by Lemmas 4 and 5, $i_{k+1} = \gamma(1-(1-\varepsilon)\mu^C)\mu_k$ and $\tau_{k+1} = \gamma(1-(1-\varepsilon)\mu^C)\mu_{k+1}$ for $0 \le k \le K-1$. Thus, Equation 7 implies $\mu_{k+1} = \beta(\gamma,\varepsilon,\mu^C)\mu_k$ for $0 \le k \le K-1$. By induction, $\mu_k = \beta(\gamma,\varepsilon,\mu^C)^k(1-\beta(\gamma,\varepsilon,\mu^C))$ for $0 \le k \le K-1$.

By Lemmas 4 and 5, $i_K = \gamma(1-(1-\varepsilon)\mu^C)\mu_{K-1}$ and $\tau_K = \gamma\mu_K$, so Equation 7 implies $\mu_K = \gamma(1-(1-\varepsilon)\mu^C)\mu_{K-1} = \beta(\gamma,\varepsilon,\mu^C)^K(1-\gamma)$. Likewise, by Lemmas 4 and 5, $i_{k+1} = \gamma\mu_k$ and $\tau_{k+1} = \gamma\mu_{k+1}$ for $k \ge K$. Hence, Equation 7 implies $\mu_{k+1} = \gamma\mu_k$ for $k \ge K$. Combining this with the previously derived $\mu_K = \beta(\gamma,\varepsilon,\mu^C)^K(1-\gamma)$ and applying induction gives $\mu_k = \gamma^{k-K}\beta(\gamma,\varepsilon,\mu^C)^K(1-\gamma)$ for $k \ge K$. This proves Equation 21.

To prove Equation 22, note that Equation 21 implies that

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

■

Lemma 16. *The value function of a player with record k is*

$$V_k = \begin{cases} (1-\beta(\gamma,\varepsilon,\mu^C)^{K-k})\mu^C & \text{if } 0 \le k \le K-1 \\ 0 & \text{if } k \ge K \end{cases} \tag{23}$$

Proof. Players with record $k \ge K$ are defectors and obtain a flow payoff of 0 in all future periods, so $V_k = 0$ for $k \ge K$. Combining this with $V_k = (1-\beta(\gamma,\varepsilon,\mu^C))\mu^C + \beta(\gamma,\varepsilon,\mu^C)V_{k+1}$ for $0 \le k \le K-1$ from Lemma 6 and solving inductively for V_k gives $V_k = (1-\beta(\gamma,\varepsilon,\mu^C)^{K-k})\mu^C$ for $0 \le k \le K-1$. ■

Lemma 17.

1. The $(C|C)_0$ constraint is

$$\frac{1-\varepsilon}{1-\gamma(1-\varepsilon)\mu^C}\beta(\gamma,\varepsilon,\mu^C)^K\mu^C > g. \tag{24}$$

2. The $(D|D)_{K-1}$ constraint is

$$\mu^C < \frac{1}{\gamma(1-\varepsilon)}\frac{l}{1+l}. \tag{25}$$

Proof. We first derive the $(C|C)_0$ constraint. From Lemma 23,

$$\begin{aligned}\gamma(1-\varepsilon)\frac{V_0-V_1}{1-\gamma} &= \gamma(1-\varepsilon)\frac{1-\beta(\gamma,\varepsilon,\mu^C)}{1-\gamma}\beta(\gamma,\varepsilon,\mu^C)^{K-1}\mu^C \\ &= \frac{1-\varepsilon}{1-(1-\varepsilon)\mu^C}\beta(\gamma,\varepsilon,\mu^C)^K\mu^C,\end{aligned}$$

and the $(C|C)_0$ constraint is equivalent to

$$\frac{1-\varepsilon}{1-\gamma(1-\varepsilon)\mu^C}\beta(\gamma,\varepsilon,\mu^C)^K\mu^C > g.$$

We now derive the $(D|D)_{K-1}$ constraint. From Lemma 23, $V_{K-1} = (1-\beta(\gamma,\varepsilon,\mu^C))\mu^C$ and $V_K = 0$. Therefore,

$$\begin{aligned}\gamma(1-\varepsilon)\frac{V_{K-1}-V_K}{1-\gamma} &= \frac{1-\beta(\gamma,\varepsilon,\mu^C)}{1-\gamma}\gamma(1-\varepsilon)\mu^C \\ &= \frac{1}{1-\gamma(1-\varepsilon)\mu^C}\gamma(1-\varepsilon)\mu^C.\end{aligned}$$

Hence, the $(D|D)_{K-1}$ constraint is equivalent to

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

Manipulating this inequality yields (25.) ■

Corollary 3. *When combined with the steady state condition Equation 22, (25) reduces* *to*

$$\mu^C \in \left(\frac{1+g-\sqrt{(1+g)^2-\frac{4}{1-\varepsilon}g}}{2}, \frac{1+g+\sqrt{(1+g)^2-\frac{4}{1-\varepsilon}g}}{2}\right). \tag{26}$$

Proof. Equation 22 implies that $\beta(\gamma,\varepsilon,\mu^C)^K = 1-\mu^C$. Combining this with Inequality 25 gives

$$(1-\mu^C)\mu^C > \frac{g}{1-\varepsilon} - g\mu^C.$$

Solving this inequality for μ^C provides the desired expression. ■

Proposition 4 follows from combining the feasibility constraint given by Equation 22 in Lemma 15 and the incentive constraints given by (25) in Lemma 17 and (26) in Corollary 3.

OA.7 Proof of Part 2 of Theorem 4

Theorem 4 (Part 2)**.** *For $g < 1$ and $l \leq g/(1-g)$,*

$$\lim_{\varepsilon\to 0}\lim_{\gamma\to 1}\overline{\mu}_K^C(\gamma,\varepsilon) = 0.$$

Proof. The case $l < g/(1-g)$ was already handled in A.6. Here we handle the case $l = g/(1-g)$, or equivalently $l/(1+l) = g$. We show that there exists some $\overline{\varepsilon} > 0$ such that $\limsup_{\gamma\to 1}\overline{\mu}_K^C(\gamma,\varepsilon) = 0$ for all $\varepsilon < \overline{\varepsilon}$. Suppose that $\limsup_{\gamma\to 1}\overline{\mu}_K^C(\gamma,\varepsilon) = \mu^C(\varepsilon) > 0$ for some ε. Then there is some $\gamma_n \to 1$ and a sequence of associated equilibria with share of cooperators $\mu^C(\gamma_n,\varepsilon)$ such that $\lim_{n\to\infty}\mu^C(\gamma_n,\varepsilon) = \mu^C(\varepsilon)$. Such a sequence must satisfy the $(C|C)_0$ and $(D|D)_{K-1}$ constraints for each corresponding γ_n. Taking the limit of these constraints as $n \to \infty$ shows that $\mu^C(\varepsilon)$ must satisfy the following "limit" constraints.

$$\text{Limit } (C|C)_0 : \mu^C(\varepsilon) \in \left[\frac{1+g-\sqrt{(1+g)^2 - \frac{4}{1-\varepsilon}g}}{2}, \frac{1+g+\sqrt{(1+g)^2 - \frac{4}{1-\varepsilon}g}}{2}\right],$$

$$\text{Limit } (D|D)_{K-1} : \mu^C(\varepsilon) \leq \frac{1}{1-\varepsilon}g.$$

We show that the function

$$q(\varepsilon) := \frac{1+g-\sqrt{(1+g)^2 - \frac{4}{1-\varepsilon}g}}{2} - \frac{1}{1-\varepsilon}g$$

is strictly positive for all sufficiently small, but non-zero, ε, which precludes $\mu^C(\varepsilon)$ satisfying the above "limit" constraints for such ε. To see that $q(\varepsilon)$ for sufficiently small, but non-zero, ε, note that $q(\varepsilon) = 0$, while the ε derivative of q evaluated at $\varepsilon = 0$ is

$$\frac{dq}{d\varepsilon}(0) = g\Big(\frac{1}{1-g} - 1\Big) > 0,$$

where the inequality comes from $0 < g < 1$. ■

OA.8 Proof of Lemma 9

Lemma 9. *Fix $\varepsilon \in (0,1)$. For all $\Delta > 0$, there exists $\overline{\gamma} < 1$ such that, for all $\gamma > \overline{\gamma}$ and $\mu^C \in [0,1]$, there exists a $\hat{\mu}^C$ satisfying $|\hat{\mu}^C - \mu^C| < \Delta$ that satisfies the Feasibility constraint of Proposition 4 for some K.*

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

$$\tilde{K}(\gamma, \varepsilon, \mu^C) = \frac{\ln(1-\mu^C)}{\ln(\beta(\gamma, \varepsilon, \mu^C))}. \tag{27}$$

By construction, $\tilde{K}(\gamma, \varepsilon, \mu^C)$ is the unique $K \in \mathbb{R}_+$ such that $\mu^C = 1 - \beta(\gamma, \varepsilon, \mu^C)^K$.

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

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

The μ^C derivative of $\tilde{K}(\gamma, \varepsilon, \mu^C)$ is related to $d(\gamma, \varepsilon, \mu^C)$ in the following lemma.

Lemma 18. $\tilde{K} : (0,1) \times (0,1) \times (0,1) \to \mathbb{R}_+$ *is differentiable in* μ^C *with derivative given by*

$$\frac{\partial \tilde{K}}{\partial \mu^C}(\gamma, \varepsilon, \mu^C) = -\frac{d(\gamma, \varepsilon, \mu^C)}{(1-\mu^C)\ln(\beta(\gamma, \varepsilon, \mu^C))}.$$

Proof of Lemma 18. From Equation 27, it follows that $\tilde{K}(\gamma, \varepsilon, \mu^C)$ is differentiable in μ^C with derivative given by

$$\begin{aligned} \frac{\partial \tilde{K}}{\partial \mu^C}(\gamma, \varepsilon, \mu^C) &= -\frac{\frac{\ln(\beta(\gamma,\varepsilon,\mu^C))}{1-\mu^C} + \frac{\ln(1-\mu^C)\frac{\partial \beta}{\partial \mu^C}(\gamma,\varepsilon,\mu^C)}{\beta(\gamma,\varepsilon,\mu^C)}}{\ln(\beta(\gamma,\varepsilon,\mu^C))^2} \\ &= -\frac{1 + \ln(1-\mu^C)(1-\mu^C)\frac{\frac{\partial \beta}{\partial \mu^C}(\gamma,\varepsilon,\mu^C)}{\beta(\gamma,\varepsilon,\mu^C)\ln(\beta(\gamma,\varepsilon,\mu^C))}}{(1-\mu^C)\ln(\beta(\gamma,\varepsilon,\mu^C))} \\ &= -\frac{d(\gamma, \varepsilon, \mu^C)}{(1-\mu^C)\ln(\beta(\gamma, \varepsilon, \mu^C))}. \end{aligned}$$ ■

The following two lemmas concern properties of $d(\gamma, \varepsilon, \mu^C)$ that will be useful for the proof of Lemma 9.

Lemma 19. $d : (0,1] \times (0,1) \times (0,1) \to \mathbb{R}$ *is well-defined and continuous.*

Proof of Lemma 19. Since $\beta(\gamma, \varepsilon, \mu^C)$ is differentiable and only takes values in $(0,1)$, it follows that $d(\gamma, \varepsilon, \mu^C)$ is well-defined. Moreover, since $\beta(\gamma, \varepsilon, \mu^C)$ is continuously differentiable for all $(\gamma, \mu^C) \in (0,1) \times (0,1)$, $d(\gamma, \varepsilon, \mu^C)$ is continuous for $\gamma < 1$. All that remains is to check that $d(\gamma, \varepsilon, \mu^C)$ is continuous for $\gamma = 1$.

First, note that $d(1, \varepsilon, \mu^C)$ is continuous in μ^C. Thus, we need only check the limit in which γ approaches 1, but never equals 1. Note that

$$\begin{aligned}\frac{\frac{\partial\beta}{\partial\mu^C}(\gamma,\varepsilon,\mu^C)}{\beta(\gamma,\varepsilon,\mu^C)\ln(\beta(\gamma,\varepsilon,\mu^C))} &= -\frac{\frac{\gamma(1-\varepsilon)(1-\gamma)}{(1-\gamma(1-\varepsilon)\mu^C)^2}}{\beta(\gamma,\varepsilon,\mu^C)\ln(\beta(\gamma,\varepsilon,\mu^C))} \\ &= -\Big(\frac{\gamma(1-\varepsilon)}{\beta(\gamma,\varepsilon,\mu^C)(1-\gamma(1-\varepsilon)\mu^C)}\Big)\Big(\frac{1-\beta(\gamma,\varepsilon,\mu^C)}{\ln(\beta(\gamma,\varepsilon,\mu^C))}\Big).\end{aligned} \tag{28}$$

It is clear that

$$\lim_{\substack{(\gamma,\mu)\to(1,\mu^C)\\ \gamma\neq 1}} \frac{\gamma(1-\varepsilon)}{\beta(\gamma,\varepsilon,\mu^C)(1-\gamma(1-\varepsilon)\mu^C)} = \frac{1-\varepsilon}{(1-(1-\varepsilon)\mu^C)} \tag{29}$$

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,

$$\lim_{\substack{(\gamma,\mu)\to(1,\mu^C)\\ \gamma\neq 1}} \frac{1-\beta(\gamma,\varepsilon,\mu^C)}{\ln(\beta(\gamma,\varepsilon,\mu^C))} = -1 \tag{30}$$

for all $\mu^C \in (0,1)$. Equations 28, 29, and 30 together imply that $d(\gamma,\varepsilon,\mu^C)$ is continuous for $\gamma = 1$. ■

Lemma 20. *For any fixed ε, $d(1,\varepsilon,\mu^C)$ has at most two zeros in $\mu^C \in (0,1)$.*

Proof of Lemma 20. It suffices to show that

$$\frac{\ln(1-\mu^C)(1-\mu^C)}{1-(1-\varepsilon)\mu^C}$$

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

$$\frac{\partial}{\partial\mu^C}\Big[\frac{\ln(1-\mu^C)(1-\mu^C)}{1-(1-\varepsilon)\mu^C}\Big] = \frac{(1-\varepsilon)\mu^C - \varepsilon\ln(1-\mu^C) - 1}{(1-(1-\varepsilon)\mu^C)^2}.$$

The single-peakedness of $\ln(1-\mu^C)(1-\mu^C)/(1-(1-\varepsilon)\mu^C)$ follows from $(1-\varepsilon)\mu^C - \varepsilon\ln(1-\mu^C) - 1$ being increasing in μ^C. ■

With these preliminaries established, we now present the proof of Lemma 9.

Proof of Lemma 9. Fix $\varepsilon \in (0,1)$. Lemma 20 says $d(1,\varepsilon,\mu^C)$ has at most two zeros for $\mu^C \in (0,1)$. 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

$$\min\{|\mu^C - \mu_1|, |\mu^C - \mu_2|, |\mu^C - \mu_3|, |\mu^C - \mu_4|, |\mu^C - \mu_5|, |\mu^C - \mu_6|\} < \Delta/2 \tag{31}$$

for all $\mu^C \in [0,1]$, and $d(1,\mu^C)$ 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 31 says that the interval endpoints can be chosen so that M is no farther than $\Delta/2$ from any $\mu^C \in [0,1]$, while the second condition implies that

$$|d(1,\mu^C)| > 0 \tag{32}$$

for all $\mu^C \in M$.

Lemma 19 says $d(\gamma,\varepsilon,\mu^C)$ is continuous for $(\gamma,\mu^C) \in (0,1] \times (0,1)$. Hence, $d(\gamma,\varepsilon,\mu^C)$ is uniformly continuous for $(\gamma,\mu^C) \in [\underline{\gamma},1] \times M$ for any $\underline{\gamma} > 0$. Equation 32 then implies that there exists some $\lambda > 0$ and $\tilde{\gamma} \in (0,1)$ such that $|d(\gamma,\varepsilon,\mu^C)/(1-\mu^C)| > \lambda$ for all $\gamma > \tilde{\gamma}$ and $\mu^C \in M$.

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 $\lim_{\gamma\to 1} \min_{\mu^C \in [0,1]} \beta(\gamma,\varepsilon,\mu^C) = 1$, there exists $\gamma' \in (0,1)$ such that $|\ln(\beta(\gamma,\varepsilon,\mu^C))| < \lambda\eta$ for all $\gamma > \gamma'$ and $\mu^C \in M$.

Moreover, $\lim_{\gamma\to 1} \min_{\mu^C \in [0,1]} \beta(\gamma,\varepsilon,\mu^C) = 1$ implies that there exists $\hat{\gamma} \in (0,1)$ such that $\tilde{K}(\gamma,\varepsilon,\mu^C) \geq 1$ for all $\gamma > \hat{\gamma}$ and $\mu^C \in M$.

Let $\overline{\gamma} = \max\{\tilde{\gamma},\gamma',\hat{\gamma}\}$. Thus, $|d(\gamma,\varepsilon,\mu^C)/((1-\mu^C)\ln(\beta(\gamma,\varepsilon,\mu^C)))| > 1/\eta$ and $\tilde{K}(\gamma,\varepsilon,\mu^C) \geq 1$ for all $\gamma > \overline{\gamma}$ and $\mu^C \in M$. For the remainder of the proof, fix $\gamma \in (\overline{\gamma},1)$. We now show that, for a given $\mu^C \in M$, there exists some $\hat{\mu}^C \in M$ and non-negative integer $\hat{K}$ such that $|\hat{\mu}^C - \mu^C| < \Delta/2$ and $\hat{\mu}^C = 1 - \beta(\gamma, \varepsilon, \hat{\mu})^{\hat{K}}$. This, when combined with Equation 31, completes the proof.

Fix $\mu^C \in M$. Suppose for concreteness that $\mu^C \in [\mu_1, \mu_2]$. An identical argument handles the case when $\mu^C \in [\mu_3, \mu_4] \cup [\mu_5, \mu_6]$. By construction, η is weakly smaller than both $(\mu_2 - \mu_1)/2$ and $\Delta/2$. Therefore, there is some $\tilde{\mu}^C \in [\mu_1, \mu_2]$ such that $\eta \leq |\tilde{\mu}^C - \mu^C| \leq \Delta/2$. Because $|d(\gamma, \varepsilon, \mu^C)/((1 - \mu^C)\ln(\beta(\gamma, \varepsilon, \mu^C)))| > 1/\eta$ for all $\mu^C \in M$, it follows from Lemma 18 that $|\partial\beta/\partial\mu^C(\gamma, \varepsilon, \mu^C)| > 1/\eta$ for all $\mu^C \in M$. Hence, $|\tilde{K}(\gamma, \varepsilon, \tilde{\mu}^C) - \tilde{K}(\gamma, \varepsilon, \mu^C)| > 1$. It thus follows that there exists some $\hat{\mu}^C$ between μ^C and $\tilde{\mu}^C$ and some non-negative integer $\hat{K}$ between $\tilde{K}(\gamma, \varepsilon, \mu^C)$ and $\tilde{K}(\gamma, \varepsilon, \tilde{\mu}^C)$ such that $\tilde{K}(\gamma, \varepsilon, \hat{\mu}^C) = \hat{K}$. Thus, $|\hat{\mu}^C - \mu^C| < \Delta/2$ and $\hat{\mu}^C = 1 - \beta(\gamma, \varepsilon, \hat{\mu}^C)^{\hat{K}}$. ■

OA.9 Proof of Proposition 5

Proposition 5. *There is a GrimKL equilibrium with total share of cooperators μ^C, share of reciprocators μ^R, and share of unconditional cooperators μ^{UC} if and only if the following conditions hold:*

1. *Feasibility:*

$$\mu^C = 1 - \alpha(\gamma, \varepsilon)^L \beta(\gamma, \varepsilon, \mu^C)^K,$$
$$\mu^R = 1 - \beta(\gamma, \varepsilon, \mu^C)^K,$$
$$\mu^{UC} = (1 - \alpha(\gamma, \varepsilon)^L)\beta(\gamma, \varepsilon, \mu^C)^K.$$

2. *Incentives:*

$$(C|C)_0 : \frac{(1-\varepsilon)(1-\mu^C)}{1-(1-\varepsilon)\mu^C}\Big[\mu^R + \mu^{UC}(l-g)\Big] > g,$$
$$(D|D)_{K-1} : \frac{\gamma(1-\varepsilon)(1-\mu^C)}{1-\gamma(1-\varepsilon)\mu^C}\Big[\mu^R + \mu^{UC}(l-g)\Big] + \mu^R l < l,$$
$$(C|D)_K : \frac{(1-\varepsilon)(1-\mu^C)}{1-(1-\varepsilon)\mu^C}\Big[\mu^R + \mu^{UC}(l-g)\Big] + \mu^R l > l.$$

Lemma 21. *In a GrimKL steady state with total share of cooperators μ^C,*

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

Moreover, μ^C satisfies the equation

$$\mu^C = 1 - \alpha(\gamma,\varepsilon)^L\beta(\gamma,\varepsilon,\mu^C)^K,$$

μ^R satisfies the equation

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

and μ^{UC} satisfies the equation

$$\mu^{UC} = (1-\alpha(\gamma,\varepsilon)^L)\beta(\gamma,\varepsilon,\mu^C)^K.$$

Proof. We establish the first part of this result. Since $i_0 = 1-\gamma$ and $\tau_0 = \gamma(1-(1-\varepsilon)\mu^C)\mu_0$, Equation 7 implies

$$\begin{aligned}\mu_0 &= \frac{1-\gamma}{1-\gamma(1-\varepsilon)\mu^C} \\ &= 1-\beta(\gamma,\varepsilon,\mu^C).\end{aligned}$$

Moreover, by Lemmas 4 and 5, $i_{k+1} = \gamma(1-(1-\varepsilon)\mu^C)\mu_k$ and $\tau_{k+1} = \gamma(1-(1-\varepsilon)\mu^C)\mu_{k+1}$ for $0 \le k \le K-1$. Thus, Equation 7 implies $\mu_{k+1} = \beta(\gamma,\varepsilon,\mu^C)\mu_k$ for $0 \le k \le K-1$. By induction, $\mu_k = \beta(\gamma,\varepsilon,\mu^C)^k(1-\beta(\gamma,\varepsilon,\mu^C))$ for $0 \le k \le K-1$.

By Lemmas 4 and 5, $i_K = \gamma(1-(1-\varepsilon)\mu^C)\mu_{K-1}$ and $\tau_K = \gamma\varepsilon\mu_K$, so Equation 7 implies

$$\mu_K = \frac{\gamma(1-(1-\varepsilon)\mu^C)}{1-\gamma(1-\varepsilon)}\mu_{K-1}$$
$$= \beta(\gamma,\varepsilon,\mu^C)^K(1-\alpha(\gamma,\varepsilon)).$$

Likewise, by Lemmas 4 and 5, $i_{k+1} = \gamma\varepsilon\mu_k$ and $\tau_{k+1} = \gamma\varepsilon\mu_{k+1}$ for $K \leq k \leq K+L-1$. Hence, Equation 7 implies $\mu_{k+1} = \alpha(\gamma,\varepsilon)\mu_k$ for $K \leq k \leq K+L-1$. Combining this with the previously derived $\mu_K = \beta(\gamma,\varepsilon,\mu^C)^K(1-\alpha(\gamma,\varepsilon))$ and applying induction gives $\mu_k = \alpha(\gamma,\varepsilon)^{k-K}\beta(\gamma,\varepsilon,\mu^C)^K(1-\alpha(\gamma,\varepsilon))$ for $K \leq k \leq K+L-1$.

By Lemmas 4 and 5, $i_{K+L} = \gamma\varepsilon\mu_{K+L-1}$ and $\tau_{K+L} = \gamma\mu_{K+L}$, so Equation 7 implies

$$\mu_{K+L} = \gamma\varepsilon\mu_{K+L-1}$$
$$= \alpha(\gamma,\varepsilon)^L\beta(\gamma,\varepsilon)^K(1-\gamma).$$

Likewise, by Lemmas 4 and 5, $i_{k+1} = \gamma\mu_k$ and $\tau_{k+1} = \gamma\mu_{k+1}$ for $k \geq K+L$. Hence, Equation 7 implies $\mu_{k+1} = \gamma\mu_k$ for $k \geq K+L$. Combining this with the previously derived $\mu_{K+L} = \alpha(\gamma,\varepsilon)^L\beta(\gamma,\varepsilon,\mu^C)^K(1-\gamma)$ and applying induction gives $\mu_k = \gamma^{k-K-L}\alpha(\gamma,\varepsilon)^L\beta(\gamma,\varepsilon,\mu^C)^K(1-\gamma)$ for $k \geq K+L$.

Now we establish the second part of the result. Using Equation 21, it follows that

$$\mu^R = \sum_{k=0}^{K-1}\mu_k = \sum_{k=0}^{K-1}\beta(\gamma,\varepsilon,\mu^C)^k(1-\beta(\gamma,\varepsilon,\mu^C))$$
$$= 1-\beta(\gamma,\varepsilon,\mu^C)^K,$$

$$\mu^{UC} = \sum_{k=K}^{K+L-1}\mu_k = \sum_{k=K}^{K+L-1}\alpha(\gamma,\varepsilon)^{k-K}\beta(\gamma,\varepsilon,\mu^C)^K(1-\alpha(\gamma,\varepsilon))$$
$$= (1-\alpha(\gamma,\varepsilon)^L)\beta(\gamma,\varepsilon,\mu^C)^K,$$

and

$$\mu^C = \mu^R + \mu^{UC} = 1-\alpha(\gamma,\varepsilon)^L\beta(\gamma,\varepsilon,\mu^C)^K,$$

which establishes Equation 22. ∎

Lemma 22. *The value function of a player with record k is*

$$V_k = \begin{cases} (1-\beta(\gamma,\varepsilon,\mu^C)^{K-k})\mu^C + \beta(\gamma,\varepsilon,\mu^C)^{K-k}(1-\alpha(\gamma,\varepsilon))^L(\mu^C - \mu^D l) & \\ +\beta(\gamma,\varepsilon,\mu^C)^{K-k}\alpha(\gamma,\varepsilon)^L\mu^{UC}(1+g) & \text{if } 0 \leq k \leq K-1 \\ (1-\alpha(\gamma,\varepsilon)^{K+L-k})(\mu^C - \mu^D l) + \alpha(\gamma,\varepsilon)^{K+L-k}\mu^{UC}(1+g) & \text{if } K \leq k \leq K+L-1 \\ \mu^{UC}(1+g) & \text{if } k \geq K+L \end{cases}$$

Proof. Players with record $k \geq K+L$ are defectors and obtain a flow payoff of $\mu^{UC}(1+g)$ in all future periods, so $V_k = \mu^{UC}(1+g)$ for $k \geq K+L$. Combining this with $V_k = (1-\alpha(\gamma,\varepsilon))(\mu^C - \mu^D l) + \alpha(\gamma,\varepsilon)V_{k+1}$ for $0 \leq k \leq K+L-1$ from Lemma 6 and solving inductively for V_k gives

$$V_k = (1-\alpha(\gamma,\varepsilon)^{K+L-k})(\mu^C - \mu^D l) + \alpha(\gamma,\varepsilon)^{K+L-k}\mu^{UC}(1+g)$$

for $K \leq k \leq K+L-1$. Finally, combining this with $V_k = (1-\beta(\gamma,\varepsilon,\mu^C))\mu^C + \beta(\gamma,\varepsilon,\mu^C)V_{k+1}$ for $0 \leq k \leq K-1$ from Lemma 6 and solving inductively for V_k gives

$$\begin{aligned} V_k = {} & (1-\beta(\gamma,\varepsilon,\mu^C)^{K-k})\mu^C + \beta(\gamma,\varepsilon,\mu^C)^{K-k}(1-\alpha(\gamma,\varepsilon)^L)(\mu^C - \mu^D l) \\ & + \beta(\gamma,\varepsilon,\mu^C)^{K-k}\alpha(\gamma,\varepsilon)^L\mu^{UC}(1+g) \end{aligned}$$

for $K \leq k \leq K+L-1$. ∎

Lemma 23.

1. *The $(C|C)_0$ constraint is*

$$\frac{1-\varepsilon}{1-(1-\varepsilon)\mu^C}\beta(\gamma,\varepsilon,\mu^C)^K\left[\mu^C - (1-\alpha(\gamma,\varepsilon)^L)(\mu^C - \mu^D l) - \alpha(\gamma,\varepsilon)\mu^{UC}(1+g)\right] > g.$$

2. The $(D|D)_{K-1}$ constraint is

$$\frac{\gamma(1-\varepsilon)}{1-\gamma(1-\varepsilon)\mu^C}\Big[\mu^C-(1-\alpha(\gamma,\varepsilon)^L)(\mu^C-\mu^D l)-\alpha(\gamma,\varepsilon)^L\mu^{UC}(1+g)\Big]<l.$$

3. The $(C|D)_K$ constraint is

$$\frac{1-\varepsilon}{\varepsilon}\alpha(\gamma,\varepsilon)^L\Big[\mu^C-\mu^D l-\mu^{UC}(1+g)\Big]>l.$$

Proof. We first derive the $(C|C)_0$ constraint. From Lemma 22,

$$V_0-V_1=(1-\beta(\gamma,\varepsilon,\mu^C))\beta(\gamma,\varepsilon,\mu^C)^{K-1}\Big[\mu^C-(1-\alpha(\gamma,\varepsilon)^L)(\mu^C-\mu^D l)-\alpha(\gamma,\varepsilon)^L\mu^{UC}(1+g)\Big].$$

Therefore,

$$\begin{aligned}\gamma(1-\varepsilon)\frac{V_0-V_1}{1-\gamma}&=\gamma(1-\varepsilon)\frac{1-\beta(\gamma,\varepsilon,\mu^C)}{1-\gamma}\beta(\gamma,\varepsilon,\mu^C)^{K-1}\Big[\mu^C-(1-\alpha(\gamma,\varepsilon)^L)(\mu^C-\mu^D l)\\&-\alpha(\gamma,\varepsilon)^L\mu^{UC}(1+g)\Big]\\&=(1-\varepsilon)\frac{\gamma}{1-\gamma(1-\varepsilon)\mu^C}\beta(\gamma,\varepsilon,\mu^C)^{K-1}\Big[\mu^C-(1-\alpha(\gamma,\varepsilon)^L)(\mu^C-\mu^D l)\\&-\alpha(\gamma,\varepsilon)^L\mu^{UC}(1+g)\Big]\\&=\frac{1-\varepsilon}{1-(1-\varepsilon)\mu^C}\beta(\gamma,\varepsilon,\mu^C)^K\Big[\mu^C-(1-\alpha(\gamma,\varepsilon)^L)(\mu^C-\mu^D l)-\alpha(\gamma,\varepsilon)^L\mu^{UC}(1+g)\Big].\end{aligned}$$

Hence, the $(C|C)_0$ constraint is equivalent to

$$\frac{1-\varepsilon}{1-(1-\varepsilon)\mu^C}\beta(\gamma,\varepsilon,\mu^C)^K\Big[\mu^C-(1-\alpha(\gamma,\varepsilon)^L)(\mu^C-\mu^D l)-\alpha(\gamma,\varepsilon)\mu^{UC}(1+g)\Big]>g.$$

We now derive the $(D|D)_{K-1}$ constraint. From Lemma 22,

$$V_{K-1}-V_K=(1-\beta(\gamma,\varepsilon,\mu^C))\Big[\mu^C-(1-\alpha(\gamma,\varepsilon)^L)(\mu^C-\mu^D l)-\alpha(\gamma,\varepsilon)^L\mu^{UC}(1+g)\Big].$$

Therefore,

$$\begin{aligned}\gamma(1-\varepsilon)\frac{V_{K-1}-V_K}{1-\gamma} &= \gamma(1-\varepsilon)\frac{1-\beta(\gamma,\varepsilon,\mu^C)}{1-\gamma}\Big[\mu^C-(1-\alpha(\gamma,\varepsilon)^L)(\mu^C-\mu^D l)-\alpha(\gamma,\varepsilon)^L\mu^{UC}(1+g)\Big] \\ &= \frac{\gamma(1-\varepsilon)}{1-\gamma(1-\varepsilon)\mu^C}\Big[\mu^C-(1-\alpha(\gamma,\varepsilon)^L)(\mu^C-\mu^D l)-\alpha(\gamma,\varepsilon)^L\mu^{UC}(1+g)\Big].\end{aligned}$$

Hence, the $(D|D)_{K-1}$ constraint is equivalent to

$$\frac{\gamma(1-\varepsilon)}{1-\gamma(1-\varepsilon)\mu^C}\Big[\mu^C-(1-\alpha(\gamma,\varepsilon)^L)(\mu^C-\mu^D l)-\alpha(\gamma,\varepsilon)^L\mu^{UC}(1+g)\Big] < l.$$

We now derive the $(C|D)_K$ constraint. From Lemma 22,

$$V_K - V_{K+1} = (1-\alpha(\gamma,\varepsilon))\alpha(\gamma,\varepsilon)^{L-1}\Big[\mu^C-\mu^D l-\mu^{UC}(1+g)\Big].$$

Therefore,

$$\begin{aligned}\gamma(1-\varepsilon)\frac{V_K-V_{K+1}}{1-\gamma} &= \gamma(1-\varepsilon)\frac{1-\alpha(\gamma,\varepsilon)}{1-\gamma}\alpha(\gamma,\varepsilon)^{L-1}\Big[\mu^C-\mu^D l-\mu^{UC}(1+g)\Big] \\ &= (1-\varepsilon)\frac{\gamma}{1-\gamma(1-\varepsilon)}\alpha(\gamma,\varepsilon)^{L-1}\Big[\mu^C-\mu^D l-\mu^{UC}(1+g)\Big] \\ &= \frac{1-\varepsilon}{\varepsilon}\alpha(\gamma,\varepsilon)^{L}\Big[\mu^C-\mu^D l-\mu^{UC}(1+g)\Big].\end{aligned}$$

Hence, the $(C|D)_K$ constraint is equivalent to

$$\frac{1-\varepsilon}{\varepsilon}\alpha(\gamma,\varepsilon)^{L}\Big[\mu^C-\mu^D l-\mu^{UC}(1+g)\Big] > l.$$

■

Corollary 4. *When combined with the steady state conditions from Lemma 21,*

1. *The $(C|C)_0$ constraint reduces to*

$$\frac{(1-\varepsilon)(1-\mu^C)}{1-(1-\varepsilon)\mu^C}\Big[\mu^R+\mu^{UC}(l-g)\Big] > g.$$

2. The $(D|D)_{K-1}$ constraint reduces to

$$\frac{\gamma(1-\varepsilon)(1-\mu^C)}{1-\gamma(1-\varepsilon)\mu^C}\left[\mu^R + \mu^{UC}(l-g)\right] + \mu^R l < l.$$

3. The $(C|D)_K$ constraint reduces to

$$\frac{(1-\varepsilon)(1-\mu^C)}{1-(1-\varepsilon)\mu^C}\left[\mu^R + \mu^{UC}(l-g)\right] + \mu^R l > l.$$

Proof. The steady state condition from Lemma 21 implies that $\beta(\gamma,\varepsilon,\mu^C)^K = 1-\mu^R$, $1-\alpha(\gamma,\varepsilon)^L = \mu^{UC}/(1-\mu^R)$, and $\alpha(\gamma,\varepsilon)^L = (1-\mu^C)/(1-\mu^R)$. Imposing these conditions on the $(C|C)_0$, $(D|D)_{K-1}$, and $(C|D)_K$ constraints in Lemma 23 and manipulating the various inequalities gives the inequalities in Corollary 5. ■

OA.10 Supporting Results for Part 1 of Theorem 5

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

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

Equation 14 can be equivalently written as

$$\rho(1,\varepsilon,\mu^C)\left[(l-g)\mu^C + (1+g-l)\mu^R\right] + l\mu^R = l.$$

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

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

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

Corollary 5. *For every $\mu^R \in (g/(1+g), 1-g/l]$,*

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

OA.10.1 Proof of Lemma 10

Lemma 10. *Fix $\mu^R \in (g/(1+g), 1-g/l]$. If $|1+\kappa(\mu^R)| > \iota(\mu^R)$, then there exists some $\overline{\varepsilon} > 0$, such that $\liminf_{\gamma \to 1} \overline{\mu}^C_{KL}(\gamma, \varepsilon) \geq h(\varepsilon, \mu^R)$ for $\varepsilon < \overline{\varepsilon}$.*

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

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

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

Lemma 24. *Fix $\mu^R \in (g/(1+g), 1-g/l]$. There exists $\overline{\varepsilon} > 0$ such that*

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

for all $\varepsilon < \overline{\varepsilon}$.

Proof of Lemma 24. Note that

$$\frac{\partial I}{\partial \mu^R}(1, \varepsilon, h(\varepsilon, \mu^R), \mu^R) = \rho(1, \varepsilon, h(\varepsilon, \mu^R))(1+g-l) + l > \rho(1, \varepsilon, h(\varepsilon, \mu^R))(1+g) > 0,$$

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

Moreover,

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

Since $\lim_{\varepsilon\to 0} h(\varepsilon,\mu^R) = 1$ and

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

it follows that

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

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

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

for all $\varepsilon < \overline{\varepsilon}$. ■

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

$$\tilde{K}(\gamma,\varepsilon,\mu^C,\mu^R) = \frac{\ln(1-\mu^R)}{\ln(\beta(\gamma,\varepsilon,\mu^C))}, \tag{33}$$

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

$$\tilde{L}(\gamma,\varepsilon,\mu^C,\mu^R) = \frac{\ln(1-\mu^C)-\ln(1-\mu^R)}{\ln(\alpha(\gamma,\varepsilon))}. \tag{34}$$

Note that $\tilde{L}(\gamma,\varepsilon,\mu^C,\mu^R)\geq 0$ whenever $\mu^C\geq\mu^R$, which is the case of interest. By construction, $\tilde{K}(\gamma, \varepsilon, \mu^C, \mu^R)$ and $\tilde{L}(\gamma, \varepsilon, \mu^C, \mu^R)$ are the unique $(K, L) \in \mathbb{R}_+ \times \mathbb{R}_+$ such that the feasibility constraints in Proposition 5 are satisfied.

Differentiating Equations 33 and 34 gives the following result.

Lemma 25. *$\tilde{K}(\gamma, \varepsilon, \mu^C, \mu^R)$ and $\tilde{L}(\gamma, \varepsilon, \mu^C, \mu^R)$ are differentiable in $(\mu^C, \mu^R) \in (0, 1) \times (0, 1)$ with partial derivatives*

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

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

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

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

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

The following lemma comes from direct calculation.

Lemma 26.

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

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

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

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

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

Lemma 27. *For all* $\varepsilon \in (0,1)$, $\zeta(\gamma,\varepsilon,\mu^C,\mu^R)$ *is continuous in* (γ,μ^C,μ^R).

Proof of Lemma 27. Clearly, $\zeta(\gamma,\varepsilon,\mu^C,\mu^R)$ is continuous whenever $\gamma < 1$. What remains is to show that it is continuous when $\gamma = 1$. Note that $\ln(1-\mu^R)\rho(1,\varepsilon,\mu^C)$ is continuous in (μ^C,μ^R). Thus, we need only check the limit in which γ approaches 1, but never equals 1. Recall that

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

It is clear that

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

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,

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

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

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

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

The following lemma concerns the extent to which, for small ε and fixed $\hat{\mu}^R \in (g/(1+g), 1-g/l]$, profiles (μ^C,μ^R) near $(h(\varepsilon,\hat{\mu}^R),\hat{\mu}^R)$ are close to feasible profiles. It combines Lemmas 26 and 27 with the inverse function theorem to obtain a bound on how far such (μ^C,μ^R) 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 $1+\zeta(1,\varepsilon,h(\varepsilon,\hat{\mu}^R),\hat{\mu}^R)$, which is close to $|1+\kappa(\hat{\mu}^R)|$ for small ε.

Lemma 28. *Fix $\hat{\mu}^R \in (g/(1+g), 1-g/l]$ and $\eta > 0$. If $|1+\kappa(\hat{\mu}^R)| > \lambda$ for some $\lambda > 0$, there exists some $\overline{\varepsilon} > 0$ such that, for all $\varepsilon < \overline{\varepsilon}$, there exists some $\overline{\gamma} < 1$ and an open neighborhood of $(h(\varepsilon,\hat{\mu}^R),\hat{\mu}^R)$, M, such that, for all $\gamma > \overline{\gamma}$ and $(\mu^C,\mu^R) \in M$, whenever $L = \tilde{L}(\gamma,\varepsilon,\mu^C,\mu^R)$ is an integer, there exists some feasible $\tilde{\mu}^C$ and $\tilde{\mu}^R$ such that*

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

Proof of Lemma 28. We handle the case where $1+\kappa(\hat{\mu}^R) > \lambda > 0$. The case where $1+\kappa(\hat{\mu}^R) < -\lambda < 0$ can be handled analogously.

Note that

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

Moreover,

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

by Lemma 5. Thus, when $1+\kappa(\hat{\mu}^R) > \lambda$, there exists some $\overline{\varepsilon} > 0$ such that, for all $\varepsilon < \overline{\varepsilon}$, there exists $\overline{\gamma}_1 < 1$ and an open neighborhood of $(h(\varepsilon,\hat{\mu}^R),\hat{\mu}^R)$, M_1, such that

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

for all $\gamma > \overline{\gamma}_1$ and $(\mu^C,\mu^R) \in M_1$. By Lemma 26, $J(\gamma,\varepsilon,\mu^C,\mu^R)$ is invertible for all such points. Thus, for a given $\varepsilon < \overline{\varepsilon}$ and $\gamma > \overline{\gamma}_1$, the inverse function theorem implies the existence of differentiable functions of (K,L), $\tilde{\mu}^C$ and $\tilde{\mu}^R$, that constitute a local inverse of $\tilde{K}$ and $\tilde{L}$ for $(\mu^C,\mu^R) \in M_1$. Additionally, the partial derivatives of these functions are given by J^{-1}, so that

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

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

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

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

and

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

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

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

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

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

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

Combining the expression for the partial derivatives of $\tilde{\mu}^C$ and $\tilde{\mu}^R$ with these inequalities gives

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

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

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

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

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

Lemma 28 then follows by noting that $\tilde{\mu}^C(\gamma, \varepsilon, \lfloor \tilde{K}(\gamma, \varepsilon, \mu^C, \mu^R) \rfloor, \tilde{L}(\gamma, \varepsilon, \mu^C, \mu^R))$

and $\tilde{\mu}^R(\gamma, \varepsilon, \lfloor \tilde{K}(\gamma, \varepsilon, \mu^C, \mu^R) \rfloor, \tilde{L}(\gamma, \varepsilon, \mu^C, \mu^R))$ is feasible whenever $\tilde{L}(\gamma, \varepsilon, \mu^C, \mu^R)$ is an integer. ■

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

$$J^C_{\hat{\mu}^R,\eta,\lambda}(\gamma, \varepsilon, \mu^C, \mu^R) = I\Big(1, \varepsilon, \mu^C - \frac{1+\eta}{\lambda}(1 - h(\varepsilon, \hat{\mu}^R)) \ln(\beta(\gamma, \varepsilon, h(\varepsilon, \hat{\mu}^R))), \mu^R\Big), \quad (35)$$

and $J^D_{\hat{\mu}^R,\eta,\lambda} : [0,1] \times (0,1) \times (0,1) \times (0,1) \to \mathbb{R}$ be the function given by

$$J^D_{\hat{\mu}^R,\eta,\lambda}(\gamma, \varepsilon, \mu^C, \mu^R) = I\Big(\gamma, \varepsilon, \mu^C, \mu^R - \frac{1+\eta}{\lambda}(1 - \hat{\mu}^R) \ln(\beta(\gamma, \varepsilon, h(\varepsilon, \hat{\mu}^R)))\Big). \quad (36)$$

Combining Lemmas 24 and 28, it follows that, if $|1 + \kappa(\hat{\mu}^R)| > \lambda$, there exists some $\overline{\varepsilon} > 0$ such that, for all $\varepsilon < \overline{\varepsilon}$ and $\eta > 0$, there exists $\overline{\gamma} < 1$ and an open neighborhood of $(h(\varepsilon, \hat{\mu}^R), \hat{\mu}^R)$, M, such that, for all $\gamma > \overline{\gamma}$ and $(\mu^C, \mu^R) \in M$, whenever $L = \tilde{L}(\gamma, \varepsilon, \mu^C, \mu^R)$ is a non-negative integer, the feasible profile $(\tilde{\mu}^C, \tilde{\mu}^R)$ described in Lemma 28 is such that $I(1, \varepsilon, \tilde{\mu}^C, \tilde{\mu}^R) \geq J^C_{\hat{\mu}^R,\eta,\lambda}(\gamma, \varepsilon, \mu^C, \mu^R)$ and $I(\gamma, \varepsilon, \tilde{\mu}^C, \tilde{\mu}^R) \leq J^D_{\hat{\mu}^R,\eta,\lambda}(\gamma, \varepsilon, \mu^C, \mu^R)$.

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

Lemma 29. *Fix $\hat{\mu}^R \in (g/(1+g), 1-g/l]$. If there is some λ such that $|1 + \kappa(\hat{\mu}^R)| > \lambda > \iota(\hat{\mu}^R)$, then there exists some $\eta > 0$ and $\overline{\varepsilon} > 0$ such that, for all $\varepsilon < \overline{\varepsilon}$,*

$$0 < \frac{\partial J^C_{\hat{\mu}^R,\eta,\lambda}}{\partial \gamma}(1, \varepsilon, h(\varepsilon, \hat{\mu}^R), \hat{\mu}^R) < \frac{\partial J^D_{\hat{\mu}^R,\eta,\lambda}}{\partial \gamma}(1, \varepsilon, h(\varepsilon, \hat{\mu}^R), \hat{\mu}^R).$$

Proof of Lemma 29. Differentiating Equation 35, we find that

$$\begin{aligned}
\frac{\partial J^C_{\hat{\mu}^R,\eta,\lambda}}{\partial \gamma}(1,\varepsilon,h(\varepsilon,\hat{\mu}^R),\hat{\mu}^R) &= -\frac{1+\eta}{\lambda}\frac{1-h(\varepsilon,\hat{\mu}^R)}{1-(1-\varepsilon)h(\varepsilon,\hat{\mu}^R)}\frac{\partial I}{\partial \mu^C}(1,\varepsilon,h(\varepsilon,\hat{\mu}^R),\hat{\mu}^R) \\
&= \frac{l(1-\hat{\mu}^R)}{1-(1-\varepsilon)h(\varepsilon,\hat{\mu}^R)}\frac{1+\eta}{\lambda}\Big(\frac{\varepsilon}{1-(1-\varepsilon)h(\varepsilon,\hat{\mu}^R)} \\
&\quad - \frac{1-h(\varepsilon,\hat{\mu}^R))}{l(1-\hat{\mu}^R)}\rho(1,\varepsilon,h(\varepsilon,\hat{\mu}^R))(l-g)\Big).
\end{aligned}$$

Differentiating Equation 36, we find that

$$\begin{aligned}
\frac{\partial J^D_{\hat{\mu}^R,\eta,\lambda}}{\partial \gamma}(1,\varepsilon,h(\varepsilon,\hat{\mu}^R),\hat{\mu}^R) &= \frac{\partial I}{\partial \gamma}(1,\varepsilon,h(\varepsilon,\hat{\mu}^R),\hat{\mu}^R) - \frac{1+\eta}{\lambda}\frac{1-\hat{\mu}^R}{1-(1-\varepsilon)h(\varepsilon,\hat{\mu}^R)}\frac{\partial I}{\partial \mu^R}(1,\varepsilon,h(\varepsilon,\hat{\mu}^R),\hat{\mu}^R) \\
&= \frac{l(1-\hat{\mu}^R)}{1-(1-\varepsilon)h(\varepsilon,\hat{\mu}^R)}\Big(1-\frac{1+\eta}{\lambda}\Big(\frac{1}{l}\rho(1,\varepsilon,h(\varepsilon,\hat{\mu}^R))(1+g-l)+1\Big)\Big).
\end{aligned}$$

Note that

$$\lim_{\varepsilon\to 0}\frac{l(1-\hat{\mu}^R)}{1-(1-\varepsilon)h(\varepsilon,\hat{\mu}^R)}\frac{\partial J^C_{\hat{\mu}^R,\eta,\lambda}}{\partial \gamma}(1,\varepsilon,h(\varepsilon,\hat{\mu}^R),\hat{\mu}^R) = \frac{1+\eta}{\lambda}\frac{(1+g)\hat{\mu}^R-g}{(1+g-l)\hat{\mu}^R+l-g}$$

and

$$\lim_{\varepsilon\to 0}\frac{l(1-\hat{\mu}^R)}{1-(1-\varepsilon)h(\varepsilon,\hat{\mu}^R)}\frac{\partial J^D_{\hat{\mu}^R,\eta,\lambda}}{\partial \gamma}(1,\varepsilon,h(\varepsilon,\hat{\mu}^R),\hat{\mu}^R) = 1-\frac{1+\eta}{\lambda}\Big(\frac{1}{(1+g-l)\hat{\mu}^R+l-g}\Big).$$

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

$$1-\frac{1}{\lambda}\Big(\frac{1}{(1+g-l)\hat{\mu}^R+l-g}\Big) > \frac{1}{\lambda}\frac{(1+g)\hat{\mu}^R-g}{(1+g-l)\hat{\mu}^R+l-g} > 0,$$

so there is some $\eta > 0$ such that

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

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

$$0 < \frac{\partial J^C_{\hat{\mu}^R,\eta,\lambda}}{\partial\gamma}(1,\varepsilon,h(\varepsilon,\hat{\mu}^R),\hat{\mu}^R) < \frac{\partial J^D_{\hat{\mu}^R,\eta,\lambda}}{\partial\gamma}(1,\varepsilon,h(\varepsilon,\hat{\mu}^R),\hat{\mu}^R)$$

for all $\varepsilon < \overline{\varepsilon}$. ∎

Lemma 30. *Fix $\hat{\mu}^R \in (g/(1+g), 1-g/l]$. There exists some $\overline{\varepsilon} > 0$ such that, for all $\varepsilon < \overline{\varepsilon}$, the isocurves of $\tilde{L}(\gamma,\varepsilon,\mu^C,\mu^R)$ and $I(1,\varepsilon,\mu^C,\mu^R)$ are not tangent at $(h(\varepsilon,\hat{\mu}^R),\hat{\mu}^R)$.*

Proof of Lemma 30. By Lemma 25, we the isocurve of $\tilde{L}(\gamma,\varepsilon,\mu^C,\mu^R)$ has slope

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

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

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

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

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

Since

$$\lim_{\varepsilon\to 0}\frac{\rho(1,\varepsilon,h(\varepsilon,\hat{\mu}^R))(1+g-l)+l}{\frac{\varepsilon}{1-(1-\varepsilon)h(\varepsilon,\hat{\mu}^R)}l-(1-h(\varepsilon,\hat{\mu}^R))(1-\hat{\mu}^R)\rho(1,\varepsilon,h(\varepsilon,\hat{\mu}^R))(l-g)} = \frac{1}{(1+g)\hat{\mu}^R-g} > 1,$$

the result follows. ∎

Combining Lemmas 29 and 30 gives the following result.

Lemma 31. *Fix $\hat{\mu}^R \in (g/(1+g), 1-g/l]$. If $|1+\kappa(\hat{\mu}^R)| > \iota(\hat{\mu}^R)$, there exists some $\overline{\varepsilon} > 0$ such that, for all $\varepsilon < \overline{\varepsilon}$ and all open neighborhoods of $(h(\varepsilon, \hat{\mu}^R), \hat{\mu}^R)$, M, there exists $\overline{\gamma} < 1$ such that, for all $\gamma > \overline{\gamma}$, there is a feasible $(\mu^C, \mu^R) \in M$ that satisfies the incentive constraints.*

Proof of Lemma 31. By Lemma 29, there exists some $\overline{\gamma} < 1$, sufficiently small neighborhood of $(\mu^C, \mu^R) = (h(\varepsilon, \hat{\mu}^R), \hat{\mu}^R)$, M, and $\eta_1, \eta_2 > 0$ such that

$$0 < \frac{\partial J^C_{\hat{\mu}^R,\eta,\lambda}}{\partial \gamma}(\gamma, \varepsilon, \mu^C, \mu^R) < \eta_1 < \eta_2 < \frac{\partial J^D_{\hat{\mu}^R,\eta,\lambda}}{\partial \gamma}(\gamma, \varepsilon, \mu^C, \mu^R)$$

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

$$\begin{aligned} J^C_{\hat{\mu}^R,\eta,\lambda}(\gamma, \varepsilon, \mu^C, \mu^R) &\geq J^C_{\hat{\mu}^R,\eta,\lambda}(1, \varepsilon, \mu^C, \mu^R) - \eta_1(1-\gamma) \\ &= I(1, \varepsilon, \mu^C, \mu^R) - \eta_1(1-\gamma) \\ J^D_{\hat{\mu}^R,\eta,\lambda}(\gamma, \varepsilon, \mu^C, \mu^R) &\leq J^D_{\hat{\mu}^R,\eta,\lambda}(1, \varepsilon, \mu^C, \mu^R) - \eta_2(1-\gamma) \\ &= I(1, \varepsilon, \mu^C, \mu^R) - \eta_2(1-\gamma) \end{aligned}$$

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

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

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

it follows that, for sufficiently large γ, isocurves of the form $I(1, \varepsilon, \mu^C, \mu^R) = (\eta_1 + \eta_2)/2(1-\gamma)$ intersect M for every μ^R in an open neighborhood of $1 - g/l$. By Lemma 30, the isocurves of $I(1, \varepsilon, \mu^C, \mu^R)$ and $\tilde{L}(\gamma, \varepsilon, \mu^C, \mu^R)$ are not tangent. Because the

$\tilde{L}(\gamma, \varepsilon, \mu^C, \mu^R)$ isocurves do not depend on γ and

$$\lim_{\gamma \to 1} \ln(\alpha(\gamma, \varepsilon)) = 0,$$

it follows by Lemma 25 that there exists some $(\mu^C, \mu^R) \in M$ on the isocurve $I(1, \varepsilon, \mu^C, \mu^R) = (\eta_1 + \eta_2)/2(1 - \gamma)$ that satisfies $\mu^R \leq 1 - g/l$ and is such that $\tilde{L}(\gamma, \varepsilon, \mu^C, \mu^R)$ is a non-negative integer for sufficiently large γ. ■

Lemma 10 is an immediate consequence of Lemma 31.

OA.10.2 Proof of Lemma 11

Lemma 11. *Suppose that $l > g(g+1)$. Some $\mu^R \in (g/(1+g), 1 - g/l]$ satisfies $|1 + \kappa(\mu^R)| > \iota(\mu^R)$ if $l > \max\{g(g+1), b(g)\}$.*

Lemma 11 is a consequence of the following lemma.

Lemma 32. *Suppose $l > g(g+1)$. Some $\mu^R \in (g/(1+g), 1-g/l]$ satisfies $|1+\kappa(\mu^R)| > \iota(\mu^R)$ 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*
$$l > \frac{3e^\phi - 2 - 2g}{\phi - 1}.$$

Proof of Lemma 32. 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

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

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. ■

Applying the special case where $\phi = 1.56$ to Lemma 32 and noting that, for $\phi = 1.56$,

$$g \geq e^{\phi} - 1 \text{ or } e^{\phi} g > \frac{3e^{\phi} - 2 - 2g}{\phi - 1}$$

only when $(1+g)/|\ln(1+g) - 1| < (3e^{\phi} - 2 - 2g)/(\phi - 1)$ or $g(g+1) > (1+g)/|\ln(1+g) - 1|$ gives Lemma 11.

OA.11 Another Family of Strategies

Fix positive integers K_1, K_2, K_3, K_4 and consider the following strategy: If $0 \leq k \leq K_1 - 1$, the player is a conditional cooperator; if $k_1 \leq K \leq k_1 + k_2 - 1$, the player is a defector, if $k_1 + k_2 \leq K \leq k_1 + k_2 + k_3 - 1$, the player is a conditional cooperator, if $k_1 + k_2 + k_3 \leq K \leq k_1 + k_2 + k_3 + k_4 - 1$, the player is an unconditional cooperator, and if $k \geq K_1 + K_2 + K_3 + K_4$, the player is a defector.

Combining Equations 7, 4, and 5 and using induction shows that the steady-state record shares are

$$\mu_k = \begin{cases} \beta(\gamma,\varepsilon,\mu^C)^k(1-\beta(\gamma,\varepsilon,\mu^C)) & \text{if } 0 \le k \le K_1 - 1 \\ \gamma^{k-K_1}\beta(\gamma,\varepsilon,\mu^C)^{K_1}(1-\gamma) & \text{if } 0 \le k - K_1 \le K_2 - 1 \\ \gamma^{K_2}\beta(\gamma,\varepsilon,\mu^C)^{k-K_2}(1-\beta(\gamma,\varepsilon,\mu^C)) & \text{if } 0 \le k - K_1 - K_2 \le K_3 - 1 \\ \gamma^{K_2}\alpha(\gamma,\varepsilon)^{k-K_1-K_2-K_3}\beta(\gamma,\varepsilon,\mu^C)^{K_1+K_3}(1-\alpha(\gamma,\varepsilon)) & \text{if } 0 \le k - K_1 - K_2 - K_3 \le K_4 - 1 \\ \gamma^{k-K_1-K_3-K_4}\alpha(\gamma,\varepsilon)^{K_4}\beta(\gamma,\varepsilon,\mu^C)^{K_1+K_3} & \text{if } k \ge K_1 + K_2 + K_3 + K_4 \end{cases}.$$

Thus,

$$\begin{aligned} \mu_{CC} &= \sum_{k=0}^{k_1-1}\mu_k + \sum_{k=k_1+k_2}^{k_1+k_2+k_3-1}\mu_k \\ &= \Big[\sum_{k=0}^{k_1-1}\beta(\gamma,\epsilon,\mu_C)^k + \sum_{k=k_1+k_2}^{k_1+k_2+k_3-1}\gamma^{k_2}\beta(\gamma,\epsilon,\mu_C)^{k-k_2}\Big](1-\beta(\gamma,\varepsilon,\mu^C)) \\ &= 1-\beta(\gamma,\epsilon,\mu_C)^{k_1} + \gamma^{k_2}\beta(\gamma,\epsilon,\mu_C)^{k_1}(1-\beta(\gamma,\epsilon,\mu_C)^{k_3}) \end{aligned} \tag{37}$$

and

$$\begin{aligned} \mu_{UC} &= \sum_{k=k_1+k_2+k_3}^{k_1+k_2+k_3+k_4-1}\mu_k \\ &= \sum_{k=k_1+k_2+k_3}^{k_1+k_2+k_3+k_4-1}\gamma^{k_2}\alpha(\gamma,\epsilon)^{k-k_1-k_2-k_3}\beta(\gamma,\epsilon,\mu_C)^{k_1+k_3}(1-\alpha(\gamma,\varepsilon)) \\ &= \gamma^{k_2}(1-\alpha(\gamma,\epsilon)^{k_4})\beta(\gamma,\epsilon,\mu_C)^{k_1+k_3}. \end{aligned} \tag{38}$$

Equations 37 and 38 give

$$\begin{aligned} \mu_C &= \mu_{CC} + \mu_{UC} \\ &= 1-\beta(\gamma,\epsilon,\mu_C)^{k_1} + \gamma^{k_2}\beta(\gamma,\epsilon,\mu_C)^{k_1} - \gamma^{k_2}\alpha(\gamma,\epsilon)^{k_4}\beta(\gamma,\epsilon,\mu_C)^{k_1+k_3}. \end{aligned} \tag{39}$$

The only incentive constraints that need to be checked are $(C|C)_0$, $(D|D)_{K_1-1}$,

$(C|C)_{K_1+K_2}$, $(D|D)_{K_1+K_2+K_3-1}$, and $(C|D)_{K_1+K_2+K_3}$. By Lemma 3, these are

$$\begin{aligned}
(C|C)_0 &: \gamma(1-\varepsilon)(V_0 - V_1) > (1-\gamma)g, \\
(D|D)_{K_1-1} &: \gamma(1-\varepsilon)(V_{K_1-1} - V_{K_1}) < (1-\gamma)l, \\
(C|C)_{K_1+K_2} &: \gamma(1-\varepsilon)(V_{K_1+K_2} - V_{K_1+K_2+1}) > (1-\gamma)g, \\
(D|D)_{K_1+K_2+K_3-1} &: \gamma(1-\varepsilon)(V_{K_1+K_2+K_3-1} - V_{K_1+K_2+K_3}) < (1-\gamma)l, \\
(C|D)_{K_1+K_2+K_3} &: \gamma(1-\varepsilon)(V_{K_1+K_2+K_3} - V_{K_1+K_2+K_3+1}) > (1-\gamma)l.
\end{aligned} \tag{40}$$

To check these incentive constraints, it suffices to compute the relevant value functions by performing the following calculations sequentially:

$$\begin{aligned}
V_{K_1+K_2+K_3+K_4} &= \mu_{UC}(1+g), \\
V_{K_1+K_2+K_3+1} &= (1-\alpha(\gamma,\varepsilon)^{K_4-1})(\mu_C - (1-\mu_C)l) + \alpha(\gamma,\varepsilon)^{K_4-1} V_{K_1+K_2+K_3+K_4}, \\
V_{K_1+K_2+K_3} &= (1-\alpha(\gamma,\varepsilon)^{K_4})(\mu_C - (1-\mu_C)l) + \alpha(\gamma,\varepsilon)^{K_4} V_{K_1+K_2+K_3+K_4}, \\
V_{K_1+K_2+K_3-1} &= (1-\beta(\gamma,\varepsilon,\mu^C))\mu^C + \beta(\gamma,\varepsilon,\mu^C) V_{K_1+K_2+K_3}, \\
V_{K_1+K_2+1} &= (1-\beta(\gamma,\varepsilon,\mu_C)^{K_3-1})\mu_C + \beta(\gamma,\varepsilon,\mu_C)^{K_3-1} V_{K_1+K_2+K_3}, \\
V_{K_1+K_2} &= (1-\beta(\gamma,\varepsilon,\mu_C)^{K_3})\mu_C + \beta(\gamma,\varepsilon,\mu_C)^{K_3} V_{K_1+K_2+K_3}, \\
V_{K_1} &= (1-\gamma^{K_2})\mu_{UC}(1+g) + \gamma^{K_2} V_{K_1+K_2}, \\
V_{K_1-1} &= (1-\beta(\gamma,\varepsilon,\mu^C))\mu^C + \beta(\gamma,\varepsilon,\mu^C) V_{K_1}, \\
V_1 &= (1-\beta(\gamma,\varepsilon,\mu_C)^{K_1-1})\mu_C + \beta(\gamma,\varepsilon,\mu_C)^{K_1-1} V_{K_1}, \\
V_0 &= (1-\beta(\gamma,\varepsilon,\mu_C)^{K_1})\mu_C + \beta(\gamma,\varepsilon,\mu_C)^{K_1} V_{K_1}.
\end{aligned} \tag{41}$$

The validity of these equations comes from combining Lemma 6 with recursion and the fact that every player with record $k \geq K_1 + K_2 + K_3 + K_4$ is a defector who faces a flow payoff of $\mu^{UC}(1+g)$ in every future period.

When $g = 1.0001$, $l = 2$, $\gamma = .99999$, and $\varepsilon = .0000000001$, we numerically verified that the strategy with $K_1 = 1$, $K_2 = 2$, $K_3 = 1$, and $K_4 = 1$ has a steady state satisfying Equations 38 and 39 with $\mu^C \approx .999984$ and $\mu^{UC} \approx .378686$. With these values, we calculated the relevant value functions using Equation 41 and showed that the constraints in Inequality 40 were satisfied. Thus, this strategy has an equilibrium even when $g > 1$ and $l < g(g+1)$.

OA.12 Stochastic Transitions

OA.12.1 Stochastic *GrimK*

This subsection shows that a stochastic version of *GrimK* strategies can support full limit cooperation whenever $g < 1$.

We use the following record-keeping system: There are two possible records, 0 and 1. Newborn players have record 0. When a player with record 0 plays D, her record transitions to 1 with probability χ. When a player with record 0 plays C, her record transitions to 1 with probability $\varepsilon\chi$. Record 1 is absorbing.

We consider *Grim*1 strategies under this record-keeping system: A player plays C if and only if both she and her opponent have record 0.

Theorem 6. *Fix parameters* $(g, l, \varepsilon, \gamma)$. *There exists* $\chi \in (0,1)$ *such that* $GRIM1$ *is a strict equilibrium with steady-state cooperation share* $\mu^C > 0$ *if and only if the following conditions hold.*

1. *Feasibility:*
$$\frac{(1-\gamma)\left(1-\mu^C\right)}{\gamma\left(1-(1-\varepsilon)\mu^C\right)\mu^C} < 1.$$

2. *Incentives:*

$$(C|C)_0 : \mu^C \in \left(\frac{1+g-\sqrt{(1+g)^2-4\frac{g}{1-\varepsilon}}}{2}, \frac{1+g+\sqrt{(1+g)^2-4\frac{g}{1-\varepsilon}}}{2}\right)$$
$$(D|D)_0 : \mu^C \notin \left[\frac{1+l-\sqrt{(1+l)^2-4\frac{l}{1-\varepsilon}}}{2}, \frac{1+l+\sqrt{(1+l)^2-4\frac{l}{1-\varepsilon}}}{2}\right].$$

Moreover, letting $\bar{\mu}^C(\gamma,\varepsilon)$ be the maximal level of μ^C that can be supported for any choice of χ, the following hold:

1. *If $g < 1$, then $\lim_{\varepsilon\to 0}\lim_{\gamma\to 1}\bar{\mu}^C(\gamma,\varepsilon) = 1$.*

2. *If $g \geq 1$, then $\bar{\mu}^C(\gamma,\varepsilon) = 0$.*

Note that $g < l$ implies that any μ^C that satisfies $(D|D)_0$ also satisfies $(C|C)_0$. In particular, $(D|D)_0$ never rules out the greatest level of μ^C that satisfies $(C|C)_0$. In addition, $(C|C)_0$ and $(D|D)_0$ are independent of γ, and Feasibility is always satisfied when γ is sufficiently large. Combined with the fact that the right endpoint of the interval describing $(C|C)_0$ converges to 1 as $\gamma \to 1$ whenever $g < 1$ (and is always at least 1 whenever $g \geq 1$), these observations imply that second part of the theorem follows immediately from the first.

Proof. Let μ^C be the population share with record 0. Let V_C be the continuation value of a player with record 0. Note that the continuation value of a player with record 1 is 0. Therefore, $V_C = (1-\gamma)\mu^C + \gamma[1-\chi(1-(1-\varepsilon)\mu^C)]V_C$, which is equivalent to

$$V_C = \frac{(1-\gamma)\mu^C}{1-\gamma+\gamma\chi(1-(1-\varepsilon)\mu^C)}.$$

On the other hand, the steady-state equation for μ^C is $1-\gamma = (1-\gamma)\mu^C + \gamma\chi(1-(1-\varepsilon)\mu^C)\mu^C$, which is equivalent to

$$\mu^C = \frac{1-\gamma}{1-\gamma+\gamma\chi(1-(1-\varepsilon)\mu^C)}. \tag{42}$$

These equations imply $V_C = \left(\mu^C\right)^2$.

It will be helpful to solve (42) for χ:

$$\chi = \frac{(1-\gamma)\left(1-\mu^C\right)}{\gamma\left(1-(1-\varepsilon)\mu^C\right)\mu^C}. \tag{43}$$

Feasibility requires that this quantity is less than 1.

Next, note that $(C|C)_0$ is $\gamma(1-\varepsilon\chi)V_C > (1-\gamma)g + \gamma(1-\chi)V_C$. The above results imply that this is equivalent to

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

Comparing this with (43), we see that $(C|C)_0$ holds iff

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

Solving this gives

$$\mu^C \in \left(\frac{1+g-\sqrt{(1+g)^2-4\frac{g}{1-\varepsilon}}}{2}, \frac{1+g+\sqrt{(1+g)^2-4\frac{g}{1-\varepsilon}}}{2}\right).$$

Similarly, $(D|D)_0$ is $\gamma(1-\varepsilon\chi)V_C < (1-\gamma)l + \gamma(1-\chi)V_C$. By the preceding results, this is equivalent to the corresponding constraint in Theorem 6.

Using (43) and solving the resulting inequality, $(D|D)_0$ holds iff the constraint in Theorem 6 holds. ■

OA.12.2 Stochastic *GrimKL*

In this subsection, we show that a stochastic version of *GrimKL* strategies can support full limit cooperation whenever either $g < 1$ or $l > g(g+1)$.

We use the following record-keeping system: There are three possible records, 0, 1, and 2. Newborn players have record 0. When a player with record 0 plays D, her record transitions to 1 with probability χ_1, while her record transitions to 1 with probability $\varepsilon\chi_1$ when she plays C. When a player with record 1 plays D, her record transitions to 2 with probability χ_2, while her record transitions to 2 with probability $\varepsilon\chi_2$ when she plays C. Record 2 is absorbing.

We consider *GrimKL* strategies under this record-keeping system, with $K = 1$ and $L = 1$: Players with record 0 are reciprocators, players with record 1 are unconditional cooperators, and players with record 2 are defectors.

Theorem 7. *Fix parameters $(g, l, \varepsilon, \gamma)$. There exist $\chi_1 \in (0,1)$ and $\chi_2 \in (0,1)$ such that GrimKL with $K = 1$ and $L = 1$ is an equilibrium with steady-state cooperation shares μ^R and μ^{UC} if and only if*

1. *Feasibility:*
$$\max\left\{\frac{(1-\gamma)\left(1-\mu^R\right)}{\gamma\left(1-(1-\varepsilon)\mu^C\right)\mu^R}, \frac{(1-\gamma)\mu^D}{\gamma\varepsilon\mu^{UC}}\right\} < 1.$$

2. *Incentives:*
$$(C|C)_0 : \frac{(1-\varepsilon)(1-\mu^C)}{1-(1-\varepsilon)\mu^C}\left[\mu^R + \mu^{UC}(l-g)\right] > g,$$
$$(D|D)_0 : \frac{(1-\varepsilon)(1-\mu^C)}{1-\gamma(1-\varepsilon)\mu^C}\left[\mu^R + \mu^{UC}(l-g)\right] < l,$$
$$(C|D)_1 : \frac{(1-\varepsilon)(1-\mu^C)}{\varepsilon(1-\mu^R)}[\mu^R - g\mu^{UC} - l(1-\mu^C)] > l.$$

Moreover, letting $\bar{\mu}^C(\gamma,\varepsilon)$ be the maximal level of μ^C that can be supported for any choice of χ_1 and χ_2, the following hold:

1. *If $g \geq 1$ and $l \geq g(1+g)$, then $\bar{\mu}^C(\gamma,\varepsilon) = 0$.*

2. *If either $g < 1$ or $l > g(1+g)$, then $\lim_{\varepsilon\to 0}\lim_{\gamma\to 1}\bar{\mu}^C(\gamma,\varepsilon) = 1$.*

Proof. We first compute the value functions. For defectors, we have $V^D = (1+g)\mu^{UC}$.

For unconditional cooperators, we have $V^{UC} = (1-\gamma)(\mu^C - l\mu^D) + \gamma(1-\varepsilon\chi_2)V^{UC} + \gamma\varepsilon\chi_2 V^D$, which is equivalent to

$$V^{UC} = \frac{(1-\gamma)(\mu^C - l\mu^D) + \gamma\varepsilon\chi_2(1+g)\mu^{UC}}{1-\gamma(1-\varepsilon\chi_2)}.$$

For reciprocators, we have $V^R = (1-\gamma)\mu^C + \gamma(1-\chi_1(1-(1-\varepsilon)\mu^C))V^R + \gamma\chi_1(1-(1-\varepsilon)\mu^C)V^{UC}$, which is equivalent to

$$V^R = \frac{(1-\gamma)\mu^C + \gamma\chi_1(1-(1-\varepsilon)\mu^C)V^{UC}}{1-\gamma+\gamma\chi_1(1-(1-\varepsilon)\mu^C)}.$$

We next consider the steady-state equations. For μ^R, we have $1-\gamma=(1-\gamma+\gamma\chi_1(1-(1-\varepsilon)\mu^C))\mu^R$, or equivalently

$$\mu^R = \frac{1-\gamma}{1-\gamma+\gamma\chi_1\left(1-(1-\varepsilon)\mu^C\right)} \iff \chi_1 = \frac{(1-\gamma)\left(1-\mu^R\right)}{\gamma\left(1-(1-\varepsilon)\mu^C\right)\mu^R}.$$

For μ^{UC}, we have

$$\gamma\chi_1\left(1-(1-\varepsilon)\mu^C\right)\mu^R = (1-\gamma+\gamma\varepsilon\chi_2)\,\mu^{UC}.$$

Using the above equation for χ_1, we can solve for χ_2 as

$$\chi_2 = \frac{(1-\gamma)\,\mu^D}{\gamma\varepsilon\mu^{UC}}.$$

Note that Feasibility says that χ_1 and χ_2 must be less than 1.

We now consider the incentive constraints. The $(C|C)_0$ constraint is $\gamma(1-\varepsilon)\chi_1(V^R-V^{UC})>(1-\gamma)g$, which is equivalent to

$$V^R - V^{UC} > \frac{(1-\gamma)\,g}{\gamma\,(1-\varepsilon)\,\chi_1}.$$

Note that

$$V^R - V^{UC} = \frac{(1-\gamma)\left(\mu^C - V^{UC}\right)}{1-\gamma+\gamma\chi_1\left(1-(1-\varepsilon)\mu^C\right)} = \mu^R\left(\mu^C - V^{UC}\right)$$

and

$$1-\gamma\,(1-\varepsilon\chi_2) = 1-\gamma+\frac{(1-\gamma)\,\mu^D}{\mu^{UC}} = (1-\gamma)\left(1+\frac{\mu^D}{\mu^{UC}}\right) = (1-\gamma)\,\frac{1-\mu^R}{\mu^{UC}}.$$

Therefore,

$$\begin{aligned}
V^{UC} &= \frac{(1-\gamma)\left(\mu^C - l\mu^D\right) + \gamma\varepsilon\chi_2 (1+g)\,\mu^{UC}}{1-\gamma(1-\varepsilon\chi_2)} = \frac{\mu^{UC}}{1-\mu^R}\left(\mu^C - l\mu^D + \frac{\mu^D}{\mu^{UC}}(1+g)\,\mu^{UC}\right) \\
&= \frac{\mu^{UC}}{1-\mu^R}\left(1-\mu^D - l\mu^D + (1+g)\,\mu^D\right) = \frac{\mu^{UC}}{1-\mu^R}\left(1-(l-g)\,\mu^D\right).
\end{aligned}$$

Thus, $(C|C)_0$ is equivalent to

$$\mu^R\left(\mu^C - \frac{\mu^{UC}}{1-\mu^R}\left(1-(l-g)\,\mu^D\right)\right) > \frac{\left(1-(1-\varepsilon)\,\mu^C\right)\mu^R g}{(1-\varepsilon)\left(1-\mu^R\right)},$$

which gives the corresponding constraint in Theorem 7.

Similarly, the $(D|D)_0$ constraint is $\gamma(1-\varepsilon)\chi_1(V^R - V^{UC}) < (1-\gamma)l$, which is equivalent to the constraint given in Theorem 7 by the previous results.

Finally, the $(C|D)_1$ constraint is $\gamma(1-\varepsilon)\chi_2(V^{UC} - V^D) > (1-\gamma)l$, which is equivalent to the constraint in Theorem 7 by the previous results.

We now consider the iterated limit. For fixed values of ε, μ^R, μ^{UC}, and μ^D, Feasibility is satisfied for high enough γ. Thus, we can ignore Feasibility and simply ask when there exist μ^R, μ^{UC}, and μ^D that satisfy the $\varepsilon \to 0$ "limit" versions of the incentive constraints:

$$\begin{aligned}
(C|C)_0 : &\qquad \mu^R + (l-g)\,\mu^{UC} > g \\
(D|D)_0 : &\qquad \mu^R + (l-g)\,\mu^{UC} < l \\
(C|D)_1 : &\quad \mu^R - g\mu^{UC} - l(1-\mu^C) > 0.
\end{aligned}$$

We show that if $g \geq 1$ and $l \leq g\,(1+g)$ then these constraints cannot be satisfied for any values of μ^R, μ^{UC}, and μ^D; while if $g < 1$ or $l > g\,(1+g)$ then they can be satisfied for values of μ^R, μ^{UC}, and μ^D such that $\mu^D = 0$. This completes the proof.

Suppose $g \geq 1$ and $l \leq g\,(1+g)$. Note that a necessary condition for $(C|D)_1$ is $\mu^R \geq \frac{g}{1+g}$; otherwise, the left-hand side of $(C|D)_1$ must be negative. Now, if $\mu^R \geq \frac{g}{1+g}$

and $g \geq 1$, for $(C|C)_0$ to hold it must be that

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

Hence, if $l < g\,(1+g)$ the constraints cannot be satisfied.

Now suppose $g < 1$ or $l > g\,(1+g)$. If $\mu^C = 1$ then $\mu^R \geq \frac{g}{1+g}$ is a sufficient condition for $(D|D)_0$. We can therefore support an equilibrium with $\mu^C = 1$ iff there exists $\mu^R \geq \frac{g}{1+g}$ such that $g < \mu^R + (l-g)\left(1-\mu^R\right) < l$, or equivalently

$$2g - l < \mu^R\,(1+g-l) < g. \tag{44}$$

Consider three cases. First, if $l = 1+g$ then $g < 1$, so $2g - l < 0 < g$ and thus (44) is trivially satisfied.

Second, if $l < 1+g$, then (44) is equivalent to

$$\frac{2g-l}{1+g-l} < \mu^R < \frac{g}{1+g-l}.$$

In this case, note that $\frac{g}{1+g} < \frac{g}{1+g-l}$, so there is a value of μ^R satisfying the constraints iff $\frac{2g-l}{1+g-l} < 1$, i.e. $g < 1$. Thus, the constraints can be satisfied if $l < 1+g$ and $g < 1$.

Third, if $l > 1+g$, then (44) is equivalent to

$$\frac{l-2g}{l-1-g} > \mu^R > -\frac{g}{l-1-g}.$$

In this case, there is a value of μ^R satisfying the constraints iff

$$\frac{g}{1+g} < \frac{l-2g}{l-1-g} \iff l > g\,(1+g)\,.$$

Thus, the constraints can be satisfied if $l > \max\{g,1\}\,(1+g)$.

Putting this together, if $g < 1$ then either $l < 1+g$ or $l > \max\{g,1\}\,(1+g)$. In either case, the constraints can be satisfied, and they can also be satisfied if $g \geq 1$ and

$l > g\,(1+g)$. ■

OA.13 Higher-Order Information

We analyze the efficiency properties of $GrimK$ when we no longer restrict the record-keeping system to only use first-order information. Here players are still reciprocators for the first K records, $0 \le k \le K-1$, and defectors for all other records, $k \ge K$, but a player has record k if the number of times she has played D and her opponent has played C is k, rather than if the number of times she has played D in total is k. As we defined the function $\beta(\gamma, \varepsilon, \mu^C)$ for the analysis of $GrimK$ when records count D's, so will it be useful here to define the function $\omega : (0,1) \times (0,1) \times [0,1] \to (0,1)$ given by

$$\omega(\gamma, \varepsilon, \mu^C) = \frac{\gamma\varepsilon(1-\varepsilon)\mu^C}{1-\gamma(1-\varepsilon(1-\varepsilon)\mu^C)}.$$

We first characterize the steady-state record shares. While $i_0 = 1-\gamma$ remains true, now $i_{k+1} = \tau_k = \gamma\varepsilon(1-\varepsilon)\mu^C\mu_k$ for all $0 \le k < K-1$, which is different than what Lemmas 4 and 5 would give. This is because a reciprocator's record only increases when she plays C and her opponent plays D. Equation 7 still applies and says that $\mu_0 = \frac{1-\gamma}{1-\gamma(1-\varepsilon)(1-\varepsilon)\mu^C} = 1-\omega(\gamma, \varepsilon, \mu^C)$ and $\mu_{k+1} = \omega(\gamma, \varepsilon, \mu^C)\mu_k$ for $0 \le k < K-1$. Induction gives $\mu_k = \omega(\gamma, \varepsilon, \mu^C)^k(1-\omega(\gamma, \varepsilon, \mu^C))$ for $0 \le k \le K-1$. Thus,

$$\mu^C = \sum_{k=0}^{K-1} \mu_k = \sum_{k=0}^{K-1} \omega(\gamma, \varepsilon, \mu^C)^k(1-\omega(\gamma, \varepsilon, \mu^C)) = 1-\omega(\gamma, \varepsilon, \mu^C)^K. \tag{45}$$

We now compute the value functions. Since defectors receive a flow payoff of 0 in every period, $V_k = 0$ for all $k \ge K$. Just as before, a reciprocator receives a flow payoff of μ^C. Moreover, before being matched in the current period, a reciprocator with record k has a probability of $1-\varepsilon(1-\varepsilon)\mu^C$ of retaining her record of k at the end of the period and a probability of $\varepsilon(1-\varepsilon)\mu^C$ of her record increasing to $k+1$. Thus,

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

for all $0 \leq k \leq K-1$, which is equivalent to

$$\begin{aligned} V_k &= \frac{1-\gamma}{1-\gamma(1-\varepsilon(1-\varepsilon)\mu^C)}\mu^C + \frac{\gamma\varepsilon(1-\varepsilon)\mu^C}{1-\gamma(1-\varepsilon(1-\varepsilon)\mu^C)}V_{k+1} \\ &= (1-\omega(\gamma,\varepsilon,\mu^C))\mu^C + \omega(\gamma,\varepsilon,\mu^C)V_{k+1}. \end{aligned}$$

Recursively solving this gives

$$V_k = (1-\omega(\gamma,\varepsilon,\mu^C)^{K-k})\mu^C \tag{46}$$

for $0 \leq k \leq K-1$.

Finally, the only incentive constraints we need worry about are the $(C|C)_k$ constraints, since a reciprocator's record never increases when the strategy calls upon her to play D. The $(C|C)_k$ constraints take the following form, which is slightly different than those in Lemma 3:

$$\gamma(1-\varepsilon(1-\varepsilon))(V_k - V_{k+1}) > (1-\gamma)g,$$

which is equivalent to

$$\gamma(1-\varepsilon(1-\varepsilon))\frac{V_k - V_{k+1}}{1-\gamma} > g.$$

By the usual argument, the $(C|C)_0$ constraint implies all other $(C|C)_k$ constraints. Furthermore, Equation 46 shows that $(C|C)_0$ is equivalent to

$$\frac{1-\varepsilon(1-\varepsilon)}{\varepsilon(1-\varepsilon)}\omega(\gamma,\varepsilon,\mu^C)^K > g.$$

Combining this with the steady-state condition given in Equation 45 implies that

$(C|C)_0$ is equivalent to

$$\mu^C < 1 - \frac{\varepsilon(1-\varepsilon)}{1-\varepsilon(1-\varepsilon)} g. \tag{47}$$

Equations 45 and 47 together give the following characterization of $GrimK$ equilibria.

Proposition 8. *There is a $GrimK$ equilibrium with total share of cooperators μ^C if and only if the following conditions hold:*

1. *Feasibility: $\mu^C = 1 - \omega(\gamma, \varepsilon, \mu^C)^K$.*

2. *Incentives: $\mu^C < 1 - \frac{\varepsilon(1-\varepsilon)}{1-\varepsilon(1-\varepsilon)} g$.*

The following result shows that $GrimK$ can always achieve limit efficiency in this setting, regardless of the values of g and l.

Theorem 8. $\lim_{\varepsilon \to 0} \lim_{\gamma \to 1} \overline{\mu}^C_K(\gamma, \varepsilon) = 1$.

Theorem 8 follows from combining $\lim_{\varepsilon \to 0} 1 - \varepsilon(1-\varepsilon)/(1-\varepsilon(1-\varepsilon))g = 1$ with the following lemma, which is an analog of Lemma 9.

Lemma 33. *Fix $\varepsilon \in (0,1)$. For all $\Delta > 0$, there exists $\overline{\gamma} < 1$ such that, for all $\gamma > \overline{\gamma}$ and $\mu^C \in [0,1]$, there exists a $\hat{\mu}^C$ satisfying $|\hat{\mu}^C - \mu^C| < \Delta$ that satisfies the Feasibility constraint of Proposition 8 for some K.*

Proof of Lemma 33. We first state the properties of the function $\omega(\gamma, \varepsilon, \mu^C)$ that we use in the proof.

1. For all $(\gamma, \varepsilon) \in (0,1) \times (0,1)$, $\omega(\gamma, \varepsilon, \mu^C)$ is continuous and non-decreasing in $\mu^C \in [0,1]$.

2. For all $(\gamma, \varepsilon) \in (0,1) \times (0,1)$, $\omega(\gamma, \varepsilon, \mu^C) > 0$ for all $\mu^C > 0$.

3. For all $(\varepsilon, \mu^C) \in (0,1) \times (0,1]$, $\lim_{\gamma \to 1} \omega(\gamma, \varepsilon, \mu^C) = 1$.

Note that, by the intermediate value theorem, for all $\gamma \in (0,1)$ and non-negative integers K, there exists some $\mu^C \in [0,1]$ such that $\mu^C = 1-\omega(\gamma, \varepsilon, \mu^C)^K$. Let $\underline{\mu}^C(\gamma, \varepsilon, K)$ denote the smallest such μ^C. That is,

$$\underline{\mu}^C(\gamma, \varepsilon, K) = \min\{\mu^C \in [0,1] : \mu^C = 1 - \omega(\gamma, \varepsilon, \mu^C)^K\}.$$

Fix $\Delta > 0$. There exists some $0 < \overline{\gamma} < 1$ such that $1-\Delta/2 < \omega(\gamma, \varepsilon, \mu^C) < 1$ for all $\gamma \in (\overline{\gamma}, 1)$ and $\mu \in [\Delta/2, 1]$. For the remainder of the proof, we assume that $\gamma \in (\overline{\gamma}, 1)$.

Since $1 - \Delta/2 < \omega(\gamma, \varepsilon, \mu^C) < 1$ for all $\mu \in [\Delta/2, 1]$, $\overline{\mu}^C(\gamma, \varepsilon, K) < \Delta/2$. Moreover, because $\lim_{K\to\infty} \overline{\mu}^C(\gamma, \varepsilon, K) = 1$, there exists some integer, $\underline{K} > 1$, such that $\overline{\mu}^C(\gamma, \varepsilon, \underline{K}) \geq 1 - \Delta/2$ and $\overline{\mu}^C(\gamma, \varepsilon, \underline{K}) < 1 - \Delta/2$ for all $K < \underline{K}$. Since $1 - \Delta/2 < \omega(\gamma, \varepsilon, \mu^C) < 1$ for all $\mu \in [\Delta/2, 1]$, it follows that both $\underline{\mu}^C(\gamma, \varepsilon, \underline{K}) < \Delta$ and $\underline{\mu}^C(\gamma, \varepsilon, K) < \underline{\mu}^C(\gamma, \varepsilon, K+1) < \underline{\mu}^C(\gamma, \varepsilon, K) + \Delta$ hold for all $K \geq \underline{K}$. These two conditions together imply that the subset $\{\underline{\mu}^C(\gamma, \varepsilon, K)\}_{K \geq \underline{K}}$ of $[0,1]$ is of distance no more than Δ from any point $\mu^C \in [0,1]$. ■