

Community Enforcement with Endogenous Records

Harry PEI*

March 14, 2024

Abstract: I study repeated games with anonymous random matching where players can erase signals from their records. When players are sufficiently long-lived and have strictly dominant actions, they will play their dominant actions with probability close to one in all equilibria. When players' expected lifespans are intermediate, there exist purifiable equilibria with a positive level of cooperation in the submodular prisoner's dilemma but not in the supermodular prisoner's dilemma. Therefore, the maximal level of cooperation a community can sustain is not monotone with respect to players' expected lifespans and the complementarity in players' actions can undermine their abilities to sustain cooperation.

Keywords: community enforcement, endogenous records, cooperation, expected lifespan, prisoner's dilemma.

*Department of Economics, Northwestern University. Email: harrydp@northwestern.edu. I thank Daron Acemoglu, Costas Cavounidis, Drew Fudenberg, Johannes Hörner, Junichiro Ishida, Michihiro Kandori, Daisuke Oyama, Satoru Takahashi, and Alexander Wolitzky for helpful comments. I thank the NSF grants SES-1947021 and SES-2337566 for financial support.

1 Introduction

When can we incentivize a group of selfish individuals to take cooperative actions? Can cooperation be sustained when people interact with different partners over time? These classic questions in economics motivated the literature on *community enforcement*. In communities with relatively few players, Kandori (1992), Ellison (1994), and Deb, Sugaya and Wolitzky (2020) show that players can cooperate even when they have no information about others' identities and histories. In communities that consist of *a large number of players* (e.g., a continuum of players), sustaining cooperation requires players to have some information about their partners' histories (Takahashi 2010). Such information is called a player's *record*, which consists of signals about his past actions and possibly also signals about his previous partners' actions.

This paper contributes to the literature on community enforcement by examining situations where records are *endogenous* in the sense that players can strategically decide whether to erase signals from their records. This is motivated by online marketplaces where participants can erase negative reviews via bribes and threats against those who left them (see Tadelis 2016 for related evidence).¹ The main takeaways are: when players can selectively include signals into their records, the maximal level of cooperation a community can sustain is not monotone with respect to the expected lifespans of its members, and that it is easier to sustain cooperation when players' actions are substitutes compared to situations where their actions are complements.

In my model, there are n populations of players. Each period, a unit mass of players are *active* in each population and all the active players are matched into groups of size n uniformly at random to play an n -player game. Each group consists of one player from each population with the player in population i playing the role of player i . Each play of the game generates a vector of signals, one for each player.

My modeling innovation is that after each player observes his realized signal, he can decide whether to include it in his *record* or to erase it. A player's *record* consists of all the signals that he has generated and included. His future partners can only observe his record but not his identity, age, the signals he erased, and the times at which the signals arrived. After each period, a certain fraction of players exit the game for exogenous reasons and are replaced by the same mass of new players with no signal in their records.

Theorem 1 shows that as long as a player has a strictly dominant action in the stage game and has a sufficiently long expected lifespan,² his dominant action will be played with probability close to 1 in all equilibria. My result applies as long as the *support* of each player's signal distribution does not depend on

¹For example, Tadelis (2016) shows that sellers' bribes and harassment against consumers who left negative reviews have caused significant biases in online ratings. Livingston (2005) finds that negative reviews are rare on eBay and sellers' sales depend mostly on the number of good reviews. Nosko and Tadelis (2015) document that only 0.07% of the reviews on eBay are negative despite their survey shows that a much larger fraction of the consumers are dissatisfied and complained to consumer service.

²In my model, a constant fraction $1 - \bar{\delta}_i$ of players from each population i exit the game after each period. Therefore, the expected lifespan of each player in population i is $(1 - \bar{\delta}_i)^{-1}$.

his partners' actions. This allows for all *first-order records* where each player's signal depends only on his own action and a large class of *second-order records*. In the repeated prisoner's dilemma, my result implies that sufficiently long-lived players will almost always defect in all equilibria. In the product choice game, my result implies that sufficiently long-lived sellers will almost never exert effort in any equilibrium.

The intuition is that when a player can erase signals, his continuation value must be non-decreasing once he includes an additional signal in his record. When this player has a strictly dominant action in the stage game, he has no incentive to take any other action unless doing so can significantly increase his continuation value. As a result, once his continuation value approaches the maximum, or equivalently, when he has a long enough good record, he will have a strict incentive to take his dominant action. When players are sufficiently long-lived, there will be a large mass of old players with long good records who have strict incentives to take their dominant actions. This will drive the average probability of cooperation down to zero.

Next, I focus on the repeated prisoner's dilemma and provide conditions under which players with *intermediate* expected lifespans can sustain cooperation in *purifiable* equilibrium (Harsanyi 1973). My motivation for the purifiability refinement is that in practice, players will always have some private payoff information and purifiable equilibria are robust to a small amount of such private information.

I show that it is impossible to sustain any cooperation in any purifiable equilibrium in the supermodular prisoner's dilemma (i.e., players have stronger incentives to cooperate when their opponents cooperate) regardless of players' patience, expected lifespans, and the monitoring structure. By contrast, in the submodular prisoner's dilemma, there exist purifiable equilibria with a positive level of cooperation when players' expected lifespans are intermediate even if only first-order records are available.

Therefore, the maximal level of cooperation a community can sustain is not monotone with respect to the expected lifespans of its members and that it is easier to sustain cooperation when players' actions are strategic substitutes compared to the case where their actions are strategic complements. These conclusions stand in contrast to ones obtained from community enforcement models where players cannot erase records, such as Takahashi (2010), Heller and Mohlin (2018), and Clark, Fudenberg and Wolitzky (2021), that it is easier to sustain cooperation when players have higher discount factors and when payoffs are supermodular.

Intuitively, due to players' ability to erase signals, those who have the highest continuation value will have a strict incentive to defect. In order to deliver those players a high continuation value, other players need to cooperate with them. When actions are complements, any player who has an incentive to cooperate with someone who always defects will have an incentive to cooperate with any other player. This contradicts the hypothesis that some of those players who always defect are supposed to receive the highest continuation value. In contrast, when actions are substitutes, players have stronger incentives to cooperate with those who

always defects. This facilitates the society to deliver high continuation values to those players.

I review the related literature in the remainder of this section. Section 2 sets up the baseline model. Section 3 shows that cooperation breaks down when players are sufficiently long-lived. Section 4 provides conditions for cooperation in the prisoner’s dilemma when players’ expected lifespans are intermediate.

Related Literature: This paper contributes to the literature on community enforcement pioneered by Kandori (1992) and Ellison (1994), and is reviewed by Wolitzky (2022). It is most closely related to community enforcement models with a large number of players (usually modeled as a continuum of players), such as Takahashi (2010), Heller and Mohlin (2018), Bhaskar and Thomas (2019), and Clark, Fudenberg and Wolitzky (2021) in which observing other players’ records is *necessary* to sustain cooperation.

Compared to those papers, I introduce the possibility of players erasing signals from their records. I show that players’ strategic manipulations of records lead to complete breakdowns in cooperation either when players have long expected lifespans or when their actions are strategic complements, but not when their expected lifespans are intermediate and their actions are substitutes. These conclusions stand in contrast to the ones in the existing literature that higher discount factors and strategic complementarities facilitate cooperation.³ My analysis also highlights the different strategic implications of players’ *time preference* and their *survival probability*, which to the best of my knowledge, is novel in the community enforcement literature. By contrast, the two are equivalent in existing models of community enforcement since the equilibrium outcomes depend only on the product of the two, i.e., the player’s discount factor.

Friedman and Resnick (2001) study repeated prisoner’s dilemma without any complementarity or substitutability in which players can erase *all* signals in their records and restart with a fresh record without any signal.⁴ In contrast, the players in my model can decide whether to erase each individual signal, which gives them more flexibility in manipulating their records. Compared to their result that cooperation is feasible when players have sufficiently high discount factors, I show that players’ ability to sustain cooperation is not monotone with respect to their discount factors. In an earlier work, Pei (2023) studies a reputation model in which a long-lived player can erase actions from their records and provides a sharp characterization of players’ equilibrium behaviors and payoffs when the cost of erasing actions is below some cutoff. Although the current paper cannot accommodate incomplete information and does not provide a full characterization

³Wiseman (2017) and Sandroni and Urgan (2018) show that higher discount factors can undermine cooperation incentives. In contrast to the current paper that focuses on repeated games, their results are obtained in *stochastic games with absorbing states*.

⁴Ghosh and Ray (1996) and Fujiwara-Greve and Okuno-Fujiwara (2009) study repeated games with voluntary separation, in which players can build long-term relationships with their current partners and after a player defects, he will be matched with a new partner from the pool of unmatched players. Their models stand in contrast to mine in which players are being rematched in every period, i.e., there is no long-term relationships between players, and moreover, each player has the flexibility to decide whether to include each new signal in their record.

of players' equilibrium behaviors, it allows for multiple long-lived players, a larger class of stage-game payoffs, and imperfect monitoring of players' actions.

This paper is also related to some recent works on dynamic information censoring such as Smirnov and Starkov (2022), Hauser (2023), and Sun (2024). Unlike those papers in which players' payoffs depend on some unknown state of the world, their payoffs depend only on the action profile in my model.

2 Model

I consider a discrete-time repeated game with a doubly infinite time horizon $k = \dots - 1, 0, 1, \dots$. There are n populations of players, indexed by $i \in I \equiv \{1, 2, \dots, n\}$. Each period, a unit mass of players are *active* from each population. By the end of period k and before the start of period $k + 1$, a fraction $1 - \bar{\delta}_i$ of the active players in population i irreversibly become inactive and are replaced by a mass $1 - \bar{\delta}_i$ of new players. Therefore, each player in population i has an expected lifespan of $(1 - \bar{\delta}_i)^{-1}$. Conditional on remaining active, each player in population i is indifferent between 1 unit of utility in period $k + 1$ and $\widehat{\delta}_i \in (0, 1)$ unit of utility in period k . Hence, they discount future payoffs by $\delta_i \equiv \widehat{\delta}_i \cdot \bar{\delta}_i$, which I call their *discount factor*. Note that players' time preference $\widehat{\delta}_i$ and their survival probability $\bar{\delta}_i$ will play different roles.

Each period, all the active players are matched into *groups* uniformly at random. Each group consists of n players, one from each population. They play an n -player finite game $\mathcal{G} \equiv \{I, A, u\}$ where $A \equiv \prod_{i=1}^n A_i$ is the set of action profiles with A_i the set of actions for the player from population i (which I refer to as *player i*), and $u_i : A \rightarrow \mathbb{R}$ is player i 's stage-game payoff. I use $a_{-i} \in A_{-i}$ to denote the actions of players other than i and I use $a \in A$ to denote an action profile. Player i maximizes his discounted average payoff $\sum_{k=1}^{+\infty} (1 - \delta_i) \delta_i^{k-1} u_i(a_{i,k}, a_{-i,k})$ with $(a_{i,k}, a_{-i,k})$ the action profile in the k th period of his life.

After taking actions, each n -player group generates a vector of signals $s \equiv (s_1, s_2, \dots, s_n)$, the distribution of which depends only on the actions of players from that group in that period. The signal for player i , denoted by s_i , is drawn from a finite set S_i according to $f_i(\cdot|a) \in \Delta(S_i)$. For every $k \geq 1$ and i , I use $s_{i,k}$ to denote player i 's signal in the k th period of his life. Without loss of generality, I assume that for every $i \in I$ and $s_i \in S_i$, there exists an action profile $a \in A$ such that $f_i(s_i|a) > 0$.

The main modeling innovation relative to the existing literature on community enforcement is that by the end of each period, each player in population i has the option to erase the realized s_i generated in that period. In another word, players can *selectively* include past signals into their records. I denote this decision by $e_{i,k} \in \{0, 1\}$ where $e_{i,k} = 1$ stands for signal $s_{i,k}$ being erased and vice versa.

For each player in population i , his *record* consists of the sequence of signals s_i that he has generated

but has not erased. Formally, the set of records for player i is denoted by $R_i \equiv \bigcup_{n=0}^{+\infty} S_i^n$. If he was born in period k , his period- k record is \emptyset and his record in period $t > k$ consists of a sequence of signals $\{s_{i,\tau_1}, \dots, s_{i,\tau_n}\}$ where $s_{i,\tau}$ is the signal generated in the τ th period of his life and $1 \leq \tau_1 < \dots < \tau_n \leq t - k$ such that $\tau \in \{\tau_1, \dots, \tau_n\}$ if and only if $e_{i,\tau} = 0$, i.e., the signal in the τ th period of his life was *not* erased.

As in Takahashi (2010) and Clark, Fudenberg and Wolitzky (2021), each player observes his own records and the *current records* of the players in his own group before choosing his action. Players *cannot* observe other players' identities, age, and the calendar times at which other players' signals were generated.

For players in population i , their *strategy* is given by $\sigma_i \equiv (\sigma_i^a, \sigma_i^e)$, where $\sigma_i^a : \prod_{i=1}^n R_i \rightarrow \Delta(A_i)$ is a mapping from his current-period record and the current-period records of other players in his own group to a distribution over his actions and $\sigma_i^e : R_i \times S_i \rightarrow [0, 1]$ is a mapping from his current-period record and the realized s_i in that period to the probability with which he erases that signal realization from his record.⁵

The solution concept is *steady state equilibrium*, or *equilibrium*, which consists of a distribution over records $\mu_i \in \Delta(R_i)$ and a strategy σ_i for each population $i \in I$ such that (i) for every $i \in I$, σ_i maximizes each player in population i 's discounted average payoff when the record distribution is $\mu \equiv \{\mu_j\}_{j=1}^n$ and the other players use strategy σ_{-i} and (ii) the record distribution is μ when players behave according to $\{\sigma_i\}_{i=1}^n$. There exists at least one equilibrium in the repeated game since there exists at least one Nash equilibrium in the stage game and all players playing the stage-game Nash equilibrium regardless of their records is an equilibrium of the repeated game. Motivated by robustness concerns, I will introduce a purifiability refinement in Section 4 when I present results on the conditions under which players can sustain cooperation.

Remarks on the Modeling Assumptions: My baseline model assumes that a player's record is a *sequence* of signals, i.e., it contains information about the *order* with which the signals arrived. All my results extend to the case where a player's record only contains information about the number of times that this player has generated each signal (i.e., the *summary statistics*), among the signals that were not erased.

In my baseline model, each player can only erase his signal in period t by the end of period t but cannot do so after period t . All my results are robust when players can also erase his period- t signal *after* period t . My results also extend to the case in which in each period, each player chooses a *subset* of all the signals that he has generated and shows this subset to his current partners before observing the subset of signals presented by his current-period partners.

⁵Although players can observe more information in addition to their records and their current partners' records (e.g., their previous partners' records, the signals they erased), focusing on this class of strategies (and equilibria resulting from these strategies) is without loss of generality in the sense that for every equilibrium in which players' strategies depend on payoff-irrelevant information, there exists an equivalent equilibrium in which all players use strategies that belong to my class and their discounted average payoffs and the average probability with which they take each action remain unchanged.

My baseline model also assumes that players incur no cost to erase signals. My results are robust when they incur an additive cost of erasing signals, as long as that cost is smaller than some cutoff (which depends only on players' stage-game payoffs).

In my model, players can erase signals but *cannot* modify their content. In practice, whether players can modify the content of signals and their costs of doing so depend on the institutional details. My assumption is motivated by the observation in Tadelis (2016) that many participants in online marketplaces post reviews because they are motivated to share their opinions, to reward other people's good behaviors and to punish bad ones, or to provide future participants useful information. If this is the case, then it seems less plausible (or at least more costly) to convince people to lie about their experiences than to ask them to stay silent. My main result, Theorem 1, shows that cooperation breaks down when players are sufficiently long-lived. This conclusion is stronger when players can only erase signals but cannot modify their content. In fact, this result also applies when players can manipulate records in other ways, as long as each of them has the option to erase their period- t signal by the end of period t .

3 Impossibility of Cooperation Between Long-Lived Players

Even when players can erase bad signals from their records, they may still be able to sustain cooperation when their discount factors are large enough. The intuition is that players can be rewarded or punished based on the *length* of their good records, which is something that they cannot manipulate.

My first theorem shows that such an intuition breaks down when players are *sufficiently long-lived*. It requires the following *non-shifting support* condition on the monitoring technology:

Condition 1. *The technology that monitors players in population i satisfies non-shifting support if for every $s_i \in S_i$, $a_i \in A_i$, and $a_{-i} \in A_{-i}$, $f_i(s_i|a_i, a_{-i}) > 0$ if and only if $f_i(s_i|a_i, a'_{-i}) > 0$ for every $a'_{-i} \in A_{-i}$.*

My non-shifting support condition only requires the *support* of the distribution over player i 's signal not depending on other players' actions. This condition can accommodate all *first-order records* where the distribution of s_i depends only on a_i . It also allows for a large class of *second-order records* as long as the support of $f_i \in \Delta(S_i)$ does not depend on a_{-i} . One example is that $s_i \equiv (s_i^1, s_i^2)$ where the distribution of s_i^1 depends only on a_i and the distribution of s_i^2 has full support under every action profile.

For any population i , given their strategy σ_i and the record distribution μ , the *average probability* with

which player i takes action $a_i \in A_i$ is the probability that the following random variable assigns to a_i :

$$\sum_{r_i \in R_i} \sum_{r_{-i} \in R_{-i}} \mu_i(r_i) \mu_{-i}(r_{-i}) \sigma_i^a(r_i, r_{-i}).$$

Theorem 1. *Suppose there exists a population $i \in \{1, 2, \dots, n\}$ whose players have a strictly dominant action a_i^* in the stage game and the technology that monitors population i satisfies non-shifting support. For every $\widehat{\delta}_i \in (0, 1)$ and $\varepsilon > 0$, there exists $\delta^* \in (0, 1)$ such that when $\bar{\delta}_i > \delta^*$, the average probability with which player i takes action a_i^* is greater than $1 - \varepsilon$ in every equilibrium.*

Theorem 1 implies that it is impossible to motivate patient players to take cooperative actions (i.e., actions other than a_i^*) when they can erase signals from their records and have sufficiently long lifespans. This conclusion does not rely on the technologies that monitor players in other populations as well as other populations' stage-game payoffs, survival probabilities, and time preferences. In the product choice game, my result implies that sufficiently long-lived sellers will almost never exert effort. In the prisoner's dilemma where all players have sufficiently long lifespans, my result implies that players will defect in almost all periods in all equilibria. As a result, their average payoffs in each equilibrium of the repeated prisoner's dilemma, which for player i is defined as

$$\sum_{r_i \in R_i} \sum_{r_{-i} \in R_{-i}} \mu_i(r_i) \mu_{-i}(r_{-i}) u_i(\sigma_i^a(r_i, r_{-i}), \sigma_{-i}^a(r_{-i}, r_i)), \quad (3.1)$$

must be arbitrarily close to their payoffs from always defecting. As will become clear later, each player's continuation value *at the time when he was born* is also close to his payoff from always defecting.

The proof is in Section 3.1. I start from a heuristic explanation. Since player i can erase signals from his records, his continuation value must be non-decreasing once he includes an additional signal in his record. When this player has a strictly dominant action a_i^* , he has an incentive to take other actions only when doing so can lead to a non-trivial increase in his continuation value. Therefore, he has no incentive to take any action other than a_i^* at records where his continuation value approaches the maximum and there is also an upper bound on the maximal length of good record that he may have in any equilibrium.

When players in population i are expected to live for a long time, there will be a large mass of old player i who have long enough good records so that their continuation values approach the maximum, as long as young player i without long good records take cooperative actions with probability bounded above 0. The average probability of cooperation in population i is low since old players have no incentive to cooperate.

3.1 Proof of Theorem 1

Fix any equilibrium (μ, σ) of the repeated game. Let $\underline{v}_i \equiv \min_{a \in A} u_i(a)$ and $\bar{v}_i \equiv \max_{a \in A} u_i(a)$ denote player i 's highest and lowest stage-game payoffs, respectively. Let

$$c^* \equiv \max_{a_i \neq a_i^*, a_{-i} \in A_{-i}} \left\{ u_i(a_i^*, a_{-i}) - u_i(a_i, a_{-i}) \right\} > 0, \quad (3.2)$$

which is player i 's lowest stage-game cost for not taking his strictly dominant action a_i^* . Let $V(r_i)$ denote player i 's continuation value when his current-period record is r_i *before* observing other players' records in his group. Let (r_i, s_i) denote player i 's record which consists of the sequence r_i followed by a signal s_i . Since player i has the option to erase s_i , record (r_i, s_i) occurs with positive probability in equilibrium only if $V(r_i, s_i) \geq V(r_i)$. This implies that $V(r_i) \geq V(\emptyset)$ for every r_i that occurs with positive probability in equilibrium. For every $k \in \mathbb{N}$, let R^k denote the set of player i 's records $r_i \in R_i$ that satisfies

$$V(r_i) \in \left[V(\emptyset) + k \frac{(1 - \delta_i)c^*}{2\delta_i}, V(\emptyset) + (k + 1) \frac{(1 - \delta_i)c^*}{2\delta_i} \right). \quad (3.3)$$

By definition, $\emptyset \in R^0$ and there exists an integer $K \leq \frac{\bar{v}_i - \underline{v}_i}{c^*} \cdot \frac{2\delta_i}{1 - \delta_i} + 1$ such that every r_i that occurs with positive probability in equilibrium belongs to $\bigcup_{j=0}^K R^j$. The rest of the proof proceeds in four steps.

Step 1: I show that there exists $D > 0$ such that for every $r_i \in R_i$ and $s_i \in S_i$ where records r_i and (r_i, s_i) occur with positive probability in equilibrium,

$$V(r_i, s_i) - V(r_i) \leq \frac{1 - \delta_i}{\delta_i} D. \quad (3.4)$$

This is the only step of the proof that uses the non-shifting support condition.

At any record $r_i \in R_i$ and for any signal $s_i \in S_i$, player i 's continuation value is no less than his discounted average payoff when he plays a_i regardless of his opponents' records and then erases the realized signal if and only if it is *not* s_i . This implies that

$$V(r_i) \geq (1 - \delta_i)\underline{v}_i + \delta_i f_i(s_i | a_i, \alpha_{-i}^*) V(r_i, s_i) + \delta_i (1 - f_i(s_i | a_i, \alpha_{-i}^*)) V(r_i),$$

where $\alpha_{-i}^* \in \Delta(A_{-i})$ is the distribution of other players' actions *conditional on* player i 's record being r_i ,

which can be computed from the equilibrium (μ, σ) via Bayes rule. The above inequality is equivalent to

$$V(r_i, s_i) \leq \frac{1 - \delta_i + \delta_i f_i(s_i|a_i, \alpha_{-i}^*)}{\delta_i f_i(s_i|a_i, \alpha_{-i}^*)} V(r_i) - \frac{(1 - \delta_i) \underline{v}_i}{\delta_i f_i(s_i|a_i, \alpha_{-i}^*)} \leq V(r_i) + \frac{1 - \delta_i}{\delta_i} \cdot \frac{\bar{v}_i - \underline{v}_i}{f_i(s_i|a_i, \alpha_{-i}^*)}. \quad (3.5)$$

Recall that for every $s_i \in S_i$, there exists $a \in A$ such that $f_i(s_i|a) > 0$. Since $f_i(\cdot|a_i, a_{-i})$ satisfies non-shifting support and S_i, A_i, A_{-i} are finite, there exists $\eta > 0$ such that for every $s_i \in S_i$, there exists $a_i \in A_i$ such that $f_i(s_i|a_i, a_{-i}) > \eta$ for every $a_{-i} \in A_{-i}$. Inequality (3.5) then implies (3.4) once we set $D \equiv \frac{\bar{v}_i - \underline{v}_i}{\eta}$.

Step 2: I derive an upper bound on the probability with which player i 's next-period record belongs to the same R^k as his current-period record, which is a function of the probability with which he takes actions other than a_i^* in that period. Let $P(R^k)$ denote the probability with which player i 's record in the next period belongs to R^k conditional on his current-period record belonging to R^k . Let $\mathcal{E}(R^k)$ denote the event that player i 's current-period record belongs to R^k and he does not take action a_i^* in the current period. Let $P'(R^k)$ denote the probability with which player i 's next period record belongs to R^k conditional on event $\mathcal{E}(R^k)$. Conditional on event $\mathcal{E}(R^k)$, player i prefers his equilibrium strategy to playing a_i^* and then erasing every realized signal that he has generated. This incentive constraint implies that

$$\mathbb{E}_{s_i} \left[\max \{V(r_i, s_i), V(r_i)\} - V(r_i) \middle| \mathcal{E}(R^k) \right] \geq \frac{1 - \delta_i}{\delta_i} c^*, \quad (3.6)$$

where c^* is defined in (3.2). Due to the upper bound on $V(r_i, s_i) - V(r_i)$ derived in (3.4) as well as the fact that $V(r_i, s_i) - V(r_i) \leq \frac{1 - \delta_i}{2\delta_i} c^*$ when (r_i, s_i) and r_i belong to the same R^k , inequality (3.6) implies that for every k such that player i takes actions other than a_i^* at R^k with strictly positive probability, we know that

$$P'(R^k) \leq 1 - \frac{c^*}{2D}. \quad (3.7)$$

Let $\pi(R^k)$ denote the probability with which player i takes actions other than a_i^* conditional on his current-period record belonging to R^k . This leads to the following upper bound on $P(R^k)$:

$$P(R^k) \leq \pi(R^k) P'(R^k) + (1 - \pi(R^k)) \leq 1 - \pi(R^k) \frac{c^*}{2D}. \quad (3.8)$$

Step 3: I derive an upper bound for $\mu_i(R^k)(1 - P(R^k))$, which is the probability of the event that player i 's current-period record belongs to R^k while his next-period record does not belong to R^k .

Let $Q(R^k \rightarrow R^j)$ denote the probability with which player i 's next-period record belongs to R^j conditional on his current-period record belonging to R^k . Since $\emptyset \in R^0$ and player i remains active in the next

period with probability $\bar{\delta}_i$, the steady state record distribution for player i , denoted by μ_i , must satisfy:

$$\mu_i(R^0) = (1 - \bar{\delta}_i) + \bar{\delta}_i \mu_i(R^0) P(R^0),$$

or equivalently,

$$\mu_i(R^0) = \frac{1 - \bar{\delta}_i}{1 - \bar{\delta}_i P(R^0)}. \quad (3.9)$$

Therefore,

$$\mu_i(R^0)(1 - P(R^0)) = \frac{(1 - \bar{\delta}_i)(1 - P(R^0))}{1 - \bar{\delta}_i P(R^0)} \leq 1 - \bar{\delta}_i. \quad (3.10)$$

For any $k \geq 1$, we have

$$\mu_i(R^k) = \bar{\delta}_i \left\{ \mu_i(R^k) P(R^k) + \sum_{j=0}^{k-1} \mu_i(R^j) Q(R^j \rightarrow R^k) \right\}$$

or equivalently,

$$\mu_i(R^k) = \frac{\bar{\delta}_i \sum_{j=0}^{k-1} \mu_i(R^j) Q(R^j \rightarrow R^k)}{1 - \bar{\delta}_i P(R^k)}. \quad (3.11)$$

This implies that

$$\mu_i(R^k)(1 - P(R^k)) = \bar{\delta}_i \frac{1 - P(R^k)}{1 - \bar{\delta}_i P(R^k)} \sum_{j=0}^{k-1} \mu_i(R^j) Q(R^j \rightarrow R^k) \leq \bar{\delta}_i \sum_{j=0}^{k-1} \mu_i(R^j) Q(R^j \rightarrow R^k). \quad (3.12)$$

Since $\sum_{j=k+1}^K Q(R^k \rightarrow R^j) = 1 - P(R^k)$ for every $k \leq K$, inequality (3.12) implies that $\mu_i(R^k)(1 - P(R^k)) \leq 2^{k-1}(1 - \bar{\delta}_i)$, and therefore,

$$\sum_{k=0}^K \mu_i(R^k)(1 - P(R^k)) \leq 2^K(1 - \bar{\delta}_i). \quad (3.13)$$

Step 4: I derive an upper bound for $\mu_i(R^k)\pi(R^k)$, which is the probability of the event that player i 's current-period record belonging to R^k but he does not take action a_i^* in the current period. This together with the upper bound on K leads to an upper bound on $\sum_{k=0}^K \mu_i(R^k)\pi(R^k)$.

Since $D \geq c^*$, equations (3.8) and (3.9) together imply that when $\bar{\delta}_i$ is close to 1,

$$\mu_i(R^0)\pi(R^0) = \frac{(1 - \bar{\delta}_i)\pi(R^0)}{1 - \bar{\delta}_i P(R^0)} \leq \frac{(1 - \bar{\delta}_i)\pi(R^0)}{1 - \bar{\delta}_i + \bar{\delta}_i \pi(R^0) \frac{c^*}{2D}} \leq (1 - \bar{\delta}_i) \cdot \frac{1}{1 - \bar{\delta}_i + \bar{\delta}_i \frac{c^*}{2D}} \leq X(1 - \bar{\delta}_i) \quad (3.14)$$

where $X \equiv \frac{2D}{c^*}$. For every $k \geq 1$, equations (3.8) and (3.11) imply that

$$\mu_i(R^k)\pi(R^k) = \frac{\bar{\delta}_i\pi(R^k)\sum_{j=0}^{k-1}\mu_1(R^j)Q(R^j \rightarrow R^k)}{1 - \bar{\delta}_iP(R^k)} \leq \bar{\delta}_iX \sum_{j=0}^{k-1}\mu_i(R^j)Q(R^j \rightarrow R^k). \quad (3.15)$$

Similar to the derivation of (3.13), we have

$$\underbrace{1 - \sum_{r \in R} \mu(r)\sigma_i^a(r)[a_i^*]}_{\text{the average probability that player } i \text{ does not play } a_i^*} = \sum_{k=0}^K \mu_i(R^k)\pi(R^k) \leq 2^K X(1 - \bar{\delta}_i). \quad (3.16)$$

Since X is a constant that depends only on stage-game payoffs and K is bounded above by a linear function of $(1 - \delta_i)^{-1}$, the right-hand-side of (3.16) converges to 0 once we fix the value of $\hat{\delta}_i$ (and hence, K is bounded from above given that $\delta_i < \hat{\delta}_i$) and send $\bar{\delta}_i \rightarrow 1$. This implies that when $\bar{\delta}_i$ is close to 1, the probability with which $\sum_{r \in R} \mu(r)\sigma_i^a(r)$ assigns to a_i^* must be close to 1.

4 Cooperation Between Players with Intermediate Lifespans

In this section, I focus on the prisoner's dilemma with two player-roles, i.e., two populations of players. I provide conditions under which players can sustain cooperation when their expected lifespans are *intermediate*. In the prisoner's dilemma, players' stage-game payoffs are

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

with $g, l > 0$.

I distinguish between the case with submodular payoff $g > l$, and the one with supermodular payoffs $g \leq l$.

Motivated by robustness concerns of mixed-strategy equilibria against players' private information about their payoffs, I provide conditions under which players can sustain cooperation in *purifiable equilibria*. The notion of purifiability was introduced by Harsanyi (1973) and was incorporated into the analysis of dynamic games by Bhaskar (1998), Bhaskar, Mailath and Morris (2013), and Bhaskar and Thomas (2019). Unlike in static games where all equilibria are purifiable under generic payoffs, there are many mixed-strategy equilibria in dynamic games that are non-purifiable (see Bhaskar 1998 for a well-known example).

My definition of purifiable equilibria in repeated games follows from that in Bhaskar and Thomas (2019). Recall that the stage game in my model is a finite normal-form game $\Gamma \equiv \{I, A, u\}$, which I refer to as the *unperturbed stage game*. For every $\varepsilon > 0$, an ε -*perturbed stage game* is denoted by $\Gamma(\varepsilon)$, in which the set

of players and the set of actions are the same as in the unperturbed stage-game but player i 's payoff from playing a_i is given by $u_i(a_i, a_{-i}) + \varepsilon z_i(a_i)$, where $z_i(a_i)$ is a random variable that is interpreted as a payoff shock that affects player i 's payoff from playing a_i . As in Harsanyi (1973), I assume that the distribution of $z_i(a_i)$ has bounded support and no atom, and that these random variables are independently distributed across actions, players, and periods. Each period, before player i acts, he observes the realizations of his own payoff shocks $\{z_i(a_i)\}_{a_i \in A_i}$ for the current period but *not* the realized payoff shocks of the other players and those in future periods. An equilibrium (σ, μ) of the unperturbed repeated game is *purifiable* if for every sequence $\varepsilon \rightarrow 0$, there exist a sequence of equilibria $(\sigma(\varepsilon), \mu(\varepsilon))$ of the ε -perturbed repeated games $\Gamma(\varepsilon)$ that converge to (σ, μ) .

I start from a result which shows that it is impossible to sustain any cooperation in any purifiable equilibrium in the *supermodular* prisoner's dilemma under any first-order record. This conclusion can be extended to general monotone-supermodular games where there exists a complete order on each player's action set under which (i) both players' payoffs have weakly increasing differences and (ii) each player's payoff is strictly increasing in their opponent's action and is strictly decreasing in their own action.

Theorem 2. *Suppose $g \leq l$. There is no purifiable equilibrium in which the average probability of playing C is strictly positive for any $\{\widehat{\delta}_i, \bar{\delta}_i\}_{i \in \{1,2\}}$ and any first-order record.*

The proof is in Section 4.1. I discuss the implications of this theorem after stating the next result.

Next, I present a positive result on the *submodular* prisoner's dilemma. In order to simplify the exposition, I focus on the case where players have the same time preference and the same expected lifespan, that is, $\widehat{\delta}_1 = \widehat{\delta}_2 \equiv \widehat{\delta}$ and $\bar{\delta}_1 = \bar{\delta}_2 \equiv \bar{\delta}$, which implies that $\delta_1 = \delta_2 \equiv \delta$. For tractability purposes, I also assume that each player's signal s_i perfectly reveals his action a_i . This allows for first-order records in which the distribution of a player's signal does not depend on other players' actions as well as any second-order record in which s_i can reveal a_i .

Theorem 3. *Suppose $g > l$ and s_i perfectly reveals a_i for every $i \in \{1,2\}$. There exists $\delta^* \in (0,1)$ such that for every $\widehat{\delta} > \delta^*$, there exists a non-empty interval $[\delta', \delta''] \subset (0,1)$ such that for every $\bar{\delta} \in [\delta', \delta'']$, there exists a purifiable equilibrium in which the average probability of playing C is strictly positive.*

The proof is in Section 4.2. The equilibrium constructed in my proof is robust to a small amount of *recording noise* considered in Clark, Fudenberg and Wolitzky (2021), that is, the space of signal realizations S_i coincides with the action space A_i and players' signals coincide with their actions with probability close to but less than 1. Theorem 3 is also robust when players in different populations have different time preferences and survival probabilities, as long as these differences are not too large.

I discuss the implications of Theorems 2 and 3. First, Theorem 3 suggests that communities that consist of members with *intermediate* expected lifespans can sustain cooperation even when these people can selectively erase records and only first-order records are available. This conclusion stands in contrast to Theorem 1 that the maximal level of cooperation goes to 0 as players' expected lifespans diverge to infinity.

The comparison between Theorems 1 and 3 suggests that the maximal level of cooperation a community can sustain is *not monotone* with respect to the expected lifespans of its members. In particular, fixing players' stage-game payoffs and time preference $\hat{\delta} \in (0, 1)$, the maximal probability with which players play C in equilibrium is strictly greater than 0 when $\bar{\delta}$ is intermediate but converges to 0 as $\bar{\delta} \rightarrow 1$. Such a non-monotonicity stands in contrast to the standard conclusions in the theory of repeated games and community enforcement that players have stronger incentives to cooperate when they have larger discount factors. The intuition is that while players with very short expected lifespans (e.g., those with $\bar{\delta}_i$ close to 0) have no incentive to cooperate due to their impatience, communities consisting of players with arbitrarily long lifespans also cannot sustain cooperation since (i) players with long good records have no incentive to cooperate when they can erase bad signals from their records and (ii) when players are sufficiently long-lived, there will be a significant mass of players with long good records as long as young players without long good records cooperate with non-trivial probability.

Second, the comparison between Theorems 2 and 3 suggests that players can sustain cooperation in purifiable equilibria in the submodular prisoner's dilemma but cannot sustain any cooperation in the supermodular prisoner's dilemma.⁶ This stands in contrast to the takeaways in Takahashi (2010), Heller and Mohlin (2018), and Clark, Fudenberg and Wolitzky (2021) that when records are first order, it is easier to sustain cooperation when payoffs are supermodular compared to the case where payoffs are submodular.⁷

For some intuition, note that a player has no incentive to cooperate when his continuation value approaches its maximum. In order to deliver a high continuation value to those players, other players need to cooperate with them with positive probability. When players' actions are strategic complements, and suppose there exists a player who has an incentive to cooperate with an opponent who always defects, he will have an incentive to cooperate with any opponent with any record. This contradicts the hypothesis that some of those players who always defect are supposed to receive the highest continuation value.

When players' actions are strategic substitutes and their expected lifespans are intermediate, they can

⁶When payoffs are supermodular, one can construct equilibria in which the average probability with which players play C is strictly positive as long as $g \leq l \leq 1 + g$ although those equilibria are not purifiable.

⁷Heller and Mohlin (2018) study a community enforcement model in which there are mixed-strategy commitment types and each player can observe k random samples of their partner's past play before choosing their actions. Takahashi (2010) and Clark, Fudenberg and Wolitzky (2021)'s negative results on the submodular prisoner's dilemma with first-order records restrict attention to grim trigger equilibria or strict equilibria, which are subsets of purifiable equilibria considered in the current paper.

sustain cooperation in the following equilibrium. Players in each population are divided into two subgroups: *juniors* with no C in their records and *seniors* with at least one C in their records. Seniors play D against all opponents. Juniors play C against seniors and mix between C and D against other juniors. Due to strategic substitutability, when a junior is indifferent between C and D against another junior who mixes between C and D , he will have a strict incentive to play C against seniors who always play D . Such an equilibrium is purifiable since all incentives are strict except for the case in which two juniors are matched.

In order for the above strategy profile to be part of an equilibrium, players' expected lifespans cannot be too short. This is because otherwise, their effective discount factors will be too low, which precludes their incentives to cooperate. Their expected lifespans cannot be too long since otherwise, there will be too many seniors in the population who are unwilling to cooperate. Anticipating this, juniors will also have no incentive to cooperate since their continuation value once they become senior is too low.

Given that players can sustain a positive level of cooperation in purifiable equilibria in the submodular prisoner's dilemma, one may wonder whether there is a folk theorem. I show that the answer is no when players can erase signals. In particular, there is a uniform upper bound $p \in (0, 1)$ on the average probability of cooperation in any equilibrium (including those that are non-purifiable) regardless of players' stage-game payoffs, time preferences, expected lifespans, and the technologies that monitor their behaviors.

Theorem 4. *There exists $p \in (0, 1)$ such that for every (l, g) , $(\widehat{\delta}_1, \widehat{\delta}_2, \bar{\delta}_1, \bar{\delta}_2)$, $\{f_i(\cdot|a_1, a_2)\}_{i \in \{1,2\}}$, and in every equilibrium, there exists a player-role whose average probability of playing C is no more than p .*

The proof is in Section 4.3. The intuition is that if the average probability with which player 1 plays C is close to 1, then with probability close to 1, player 2 will have records such that *unconditional on* player 1's record, player 1 will play C against these records of player 2's with probability close to 1. Since player 2 has the ability to erase signals, player 2 at such records can secure a payoff close to $1 + g$ by playing D in every period and erasing every realized signal. This payoff lower bound leads to an upper bound on the average probability with which player 2 plays C , which is bounded below 1.

4.1 Proof of Theorem 2

Suppose by way of contradiction that there exists a purifiable equilibrium where player 1 plays C with positive probability. Let $V_1(r)$ denote player 1's expected continuation value when his record is r , before observing his current partner's record. Let $\bar{V}_1 \equiv \sup_{r \in R} V_1(r)$ denote the supremum of player 1's continuation value within this equilibrium. Since player 1 can erase signals from his records, for every $\widehat{\delta}_1, \bar{\delta}_1 \in (0, 1)$, there exists $\varepsilon > 0$ such that player 1 strictly prefers to play D at any record r^* that satis-

fies $V_1(r^*) \geq \bar{V}_1 - \varepsilon$. Since player 1 plays C with positive probability in this equilibrium, it cannot be the case that he has the same continuation value at all of his records. Therefore, there exists r^* such that $V_1(r^*) > \bar{V}_1 - \varepsilon$ and $V_1(r^*) > V_1(\emptyset)$. Hence, there exists a record r' of player 2's such that r' plays C with strictly higher probability against player 1 with record r^* than against player 1 with record \emptyset .

When $g < l$, this is incentive compatible only when player 1 with record \emptyset plays D for sure against player 2 with record r' . As a result, player 2 with record r' plays C with different probabilities against player 1 with records r^* and \emptyset despite his payoffs are the same when he faces player 1 with record r^* and when he faces player 1 with record \emptyset . This contradicts the hypothesis that the equilibrium is purifiable.

Similarly, when $g = l$, player 2 with record r' needs to play C with different probabilities against player 1 with record r^* and those with record \emptyset despite his payoffs are the same when he faces player 1 with record r^* and when he faces player 1 with record \emptyset . Such a mixed strategy is not purifiable.

4.2 Proof of Theorem 3

I construct a purifiable equilibrium where the average probability with which players play C is strictly positive when $\bar{\delta}$ is intermediate and $g > l$. I partition the set of records into two categories. I say that a player is *junior* if there is no C in his record and is *senior* otherwise. The number of D a player has in his record does *not* affect whether he is senior or junior. Seniors play D against everyone. Juniors play C against seniors and play C with probability $q \in (0, 1)$ against other juniors. Let μ_0 and μ_1 denote the probabilities with which a player is junior and is senior. Let V_0 and V_1 denote the continuation values of a junior and of a senior. Each junior is indifferent between C and D when facing another junior, which gives:

$$(1 - \delta)u_1(C, q) + \delta V_1 = (1 - \delta)u_1(D, q) + \delta V_0,$$

or equivalently,

$$V_1 - V_0 = \frac{1 - \delta}{\delta} (qg + (1 - q)l). \quad (4.1)$$

This incentive constraint leads to a lower bound on δ since $V_0 \geq 0$ and $V_1 \leq 1 + g$.

Since payoffs are strictly submodular, a junior's incentive to play C against another junior implies his *strict* incentive to play C against any senior. Each junior's continuation value satisfies

$$V_0 = \mu_0 \left\{ (1 - \delta)u_1(C, q) + \delta V_1 \right\} + \mu_1 \left\{ (1 - \delta)u_1(C, D) + \delta V_1 \right\}$$

which from (4.1), is equivalent to

$$V_0 = q(\mu_0 + g - \mu_1 l). \quad (4.2)$$

Since a senior's continuation value is $V_1 = \mu_0(1 + g)$, equations (4.1) and (4.2) together imply that

$$\mu_0(1 + g) = q(\mu_0 + g - \mu_1 l) + \frac{1 - \delta}{\delta} (qg + (1 - q)l),$$

or equivalently,

$$(1 - q)\mu_0 = \frac{1 - \delta}{\delta} \cdot \frac{l}{1 + g} + q \frac{g - l}{1 + g} \left(\frac{1}{\delta} - \mu_0 \right). \quad (4.3)$$

The steady state record distribution (μ_0, μ_1) satisfies

$$\mu_0 = \underbrace{(1 - \bar{\delta})}_{\text{newborns are junior}} + \bar{\delta} \mu_0 \cdot \underbrace{\mu_0(1 - q)}_{\text{the probability that a junior remains junior in the next period}}, \quad (4.4)$$

which is a second order polynomial of μ_0 . For any $q \in (0, 1)$, there are two solutions to (4.4) and only the smaller one belongs to the interval $(0, 1)$. As a result, μ_0 is strictly decreasing in q , where $\mu_0 = 1 - \bar{\delta}$ when $q = 1$ and $\mu_0 = (1 - \bar{\delta})/\bar{\delta}$ when $q = 0$. Since the LHS of (4.3) is strictly decreasing in q and equals 0 when $q = 1$ and the RHS of (4.3) is strictly increasing in q and is always strictly positive, there exists a solution to (4.3) and (4.4) if and only if the LHS is no less than the RHS when $q = 0$, or equivalently,

$$\frac{1 - \bar{\delta}}{\bar{\delta}} \geq \frac{1 - \delta}{\delta} \cdot \frac{l}{1 + g}. \quad (4.5)$$

This leads to an upper bound on $\bar{\delta}$ for any fixed $\hat{\delta} \in (0, 1)$. Since $\bar{\delta} > \delta$, for every fixed $\hat{\delta}$ large enough (recall the lower bound on δ implied by (4.1)), an equilibrium with a positive level of cooperation exists when $\bar{\delta}$ belongs to some interval $[\delta', \delta''] \subset (0, 1)$. The purifiability of this equilibrium then follows from standard arguments since players have strict incentives except when two juniors are matched.

4.3 Proof of Theorem 4

Suppose by way of contradiction that there does not exist such a uniform upper bound $p \in (0, 1)$ on the average probability of playing C . Then for every $\varepsilon > 0$, there exist a parameter configuration $l, g > 0$, $(\hat{\delta}_1, \hat{\delta}_2, \bar{\delta}_1, \bar{\delta}_2) \in [0, 1]^4$, and $\{f_i(\cdot | a_1, a_2)\}_{i \in \{1, 2\}}$ as well as an equilibrium under these parameters such that each player-role's average probability of cooperation is at least $1 - \varepsilon$. Let μ_i denote the steady state distribution of player i 's records. Let $R_{i\eta}$ denote the set of player i 's records such that the probability with

which player $-i$ cooperates with these records is at least $1 - \eta$. By definition, $\mu_i(R_\eta) \geq 1 - \frac{\varepsilon}{\eta}$. Since player i can take action D and erase whatever signal he generates, his average continuation value is at least $(1+g)\left(1 - \frac{\varepsilon}{\eta}\right)(1-\eta)$. When ε is small enough, the above expression is strictly greater than $(1-\varepsilon) + \varepsilon(1+g)$, which is the maximal payoff a player can obtain when he cooperates with probability more than $1 - \varepsilon$. This contradicts the hypothesis that his average probability of cooperation being greater than $1 - \varepsilon$.

References

- [1] Bhaskar, V (1998) “Informational Constraints and the Overlapping Generations Model: Folk and Anti-Folk Theorems,” *Review of Economic Studies*, 65(1), 135-149.
- [2] Bhaskar, V., George Mailath and Stephen Morris (2013) “A Foundation for Markov Equilibria in Sequential Games with Finite Social Memory,” *Review of Economic Studies*, 80(3), 925-948.
- [3] Bhaskar, V and Caroline Thomas (2019) “Community Enforcement of Trust with Bounded Memory,” *Review of Economic Studies*, 86(3), 1010-1032.
- [4] Clark, Daniel, Drew Fudenberg and Alexander Wolitzky (2021) “Record-Keeping and Cooperation in Large Societies,” *Review of Economic Studies*, 88, 2179-2209.
- [5] Deb, Joyee, Takuo Sugaya and Alexander Wolitzky (2020) “The Folk Theorem in Repeated Games With Anonymous Random Matching,” *Econometrica*, 88(3), 917-964.
- [6] Ellison, Glenn (1994) “Cooperation in the Prisoner’s Dilemma with Anonymous Random Matching,” *Review of Economic Studies*, 61(3), 567-588.
- [7] Friedman, Eric and Paul Resnick (2001) “The Social Cost of Cheap Pseudonyms,” *Journal of Economics & Management Strategy*, 10(2), 173-199.
- [8] Fujiwara-Greve, Takako and Masahiro Okuno-Fujiwara (2009) “Voluntarily Separable Repeated Prisoner’s Dilemma,” *Review of Economic Studies*, 76(3), 993-1021.
- [9] Ghosh, Parikshit and Debraj Ray (1996) “Cooperation in Community Interaction without Information Flows,” *Review of Economic Studies*, 63(3), 491-519.
- [10] Grossman, Sanford (1981) “The Informational Role of Warranties and Private Disclosure about Product Quality,” *Journal of Law & Economics*, 24(3), 461-483.
- [11] Harsanyi, John (1973) “Games with Randomly Distributed Payoffs: A New Rationale for Mixed Strategy Equilibrium Points,” *International Journal of Game Theory*, 2, 1-23.
- [12] Hauser, Daniel (2023) “Censorship and Reputation,” *American Economic Journal: Microeconomics*, 15(1), 497-528.
- [13] Heller, Yuval and Erik Mohlin (2018) “Observations on Cooperation,” *Review of Economic Studies*, 88, 1892-1935.
- [14] Kandori, Michihiro (1992) “Social Norms and Community Enforcement,” *Review of Economic Studies*, 59(1), 63-80.

- [15] Livingston, Jeffrey (2005) “How Valuable Is a Good Reputation? A Sample Selection Model of Internet Auctions,” *Review of Economics and Statistics*, 87(3),453-465.
- [16] Milgrom, Paul (1981) “Good News and Bad News: Representation Theorems and Applications,” *Bell Journal of Economics*, 12(2), 380-391.
- [17] Nosko, Chris and Steven Tadelis (2015) “The Limits of Reputation in Platform Markets: An Empirical Analysis and Field Experiment,” NBER Working Paper, 20830.
- [18] Pei, Harry (2023) “Reputation Effects with Endogenous Records,” Working Paper.
- [19] Sandroni, Alvaro and Can Urgan (2018) “When to Confront: The Role of Patience,” *American Economic Journal: Microeconomics*, 10(3), 219-252.
- [20] Smirnov, Aleksei, and Egor Starkov (2022) “Bad News Turned Good: Reversal Under Censorship,” *American Economic Journal: Microeconomics*, 14(2), 506-560.
- [21] Sun, Yiman (2024) “A Dynamic Model of Censorship,” *Theoretical Economics*, 19(1), 29-60.
- [22] Tadelis, Steven (2016) “Reputation and Feedback Systems in Online Platform Markets,” *Annual Review of Economics*, 8, 321-340.
- [23] Takahashi, Satoru (2010) “Community Enforcement When Players Observe Partners’ Past Play,” *Journal of Economic Theory*, 145, 42-62.
- [24] Wiseman, Thomas (2017) “When Does Predation Dominate Collusion?” *Econometrica*, 85(2), 555-584.
- [25] Wolitzky, Alexander (2022) “Cooperation in Large Societies,” *Advances in Economics and Econometrics*, 12th World Congress of the Econometric Society.