Community Enforcement of Trust*

V Bhaskar† Caroline Thomas‡

January 29, 2018

Abstract

We examine how trust is sustained in large societies with random matching, when the player to be trusted may default voluntarily or involuntarily. In order to incentivize trustworthiness, defaulters should be punished through temporarily exclusion.The difficulty is that trusting defaulters who are the verge of rehabilitation is profitable. With perfect bounded information, defaulter exclusion unravels and trust cannot be sustained. A coarse information structure that pools recent defaulters with those nearing rehabilitation endogenously generates adverse selection, sustaining the temporary exclusion of defaulters. Equilibria where defaulters are trusted with positive probability improve efficiency, since mixing raises the proportion of likely re-offenders in the pool of defaulters. Our results extend to a large class of sequential-move games.

JEL codes: C73, D82, G20, L14, L15.

Keywords: trust game, moral hazard, repeated games with community enforcement, imperfect monitoring, bounded memory, information design.

---

*Thanks to Andrew Atkeson, Dirk Bergemann, Mehmet Ekmekci, Mark Feldman, John Geanakoplos, Andy Glover, Johannes Hörner, George Mailath, Larry Samuelson, Tom Wiseman and seminar audiences at Austin, Chicago, NYU, Stanford, Turin, Toronto, Toulouse, Western Ontario, and Yale for helpful comments. We are grateful to the Cowles Foundation at Yale for its hospitality while this paper was written. Bhaskar thanks the National Science Foundation for its support via grant 201503942.

†Department of Economics, University of Texas at Austin. v.bhaskar@austin.utexas.edu.

‡Department of Economics, University of Texas at Austin. caroline.thomas@austin.utexas.edu.

1 Introduction

We examine information and rating systems designed to induce cooperation, in large societies where interactions are bilateral and moral hazard is one-sided. Our leading application is the trust game, although our results extend to a large class of sequential-move games. The trust game captures many economic situations: when the buyer of a product places an order, the seller must decide how diligently to execute it; when a house-owner engages a builder to refurbish his house, the builder knows that shoddy work may temporarily go undetected; when a borrower takes an unsecured loan, she must subsequently decide whether to repay or wilfully default. The society is large and each pair of agents transacts infrequently. Thus, opportunistic behavior (by the seller, builder or borrower) can be deterred only by a "reputational mechanism", whereby opportunism results in future exclusion.

We assume that information on past transgressions is subject to *bounded social memory* and is retained only for a finite length of time. While this is plausible in any context, in many contexts it is legally mandated — notably in consumer credit markets. In the United States, the bankruptcy "flag" of an individual filing for bankruptcy under Chapter 7 remains on her record for 10 years, and must then be removed; if she files under Chapter 13, it remains on her record for 7 years. Elul and Gottardi (2015) find that, among the 113 countries with credit bureaus, 90 percent have time-limits on the reporting of adverse information concerning borrowers. Bounded memory also arises under policies used by internet platforms to compute the scores summarizing their participants' reputations. For example, Amazon lists a summary statistic of seller performance over the past 12 months — given that buyers have limited attention, this may serve as to effectively limit memory. In the United States, 24 states and many municipalities[1] have introduced "ban the box" legislation, prohibiting employers from asking job applicants about prior convictions unless those relate directly to the job.[2]

How do societies enforce trustworthiness when moral hazard is important and information systems are constrained by bounded memory? For the sake of exposition, consider the credit market interpretation of the trust game.[3] Our model has a large number of long-lived

[1]See "Pandora's box" in *The Economist*, August 13th 2016.

[2]One should note the broader philosophical appeal of the principle, implicit to bounded memory, that an individual's transgressions in the distant past should not be perpetually held against them. This is embodied in the European Court of Justice's determination that individuals have the "right to be forgotten", and may therefore compel online search engines to delete past records pertaining to them.

[3]This paper does not provide an all-encompassing model of consumer credit markets — our focus is on the community enforcement of trust in large random matching environments.

borrowers and lenders, where each borrower-lender pair interacts only once.[4] Lending is efficient and profitable for the lender, provided that the borrower intends to repay the loan. However, the borrower is subject to moral hazard, and has short-term incentives to wilfully default. Additionally, there is a small chance of involuntary default. Thus, lending can only be supported via long-term repayment incentives whereby default results in the borrower's future exclusion from credit.

In our large-population random-matching environment, each lender is only concerned with the profitability of his current loan. As long as he expects that loan to be repaid, he has no interest in punishing a borrower for her past transgressions. Thus, a borrower can only be deterred from wilful default if a defaulter's record indicates that she is likely to default on a subsequent loan.[5] With bounded memory, *disciplining lenders* to not lend to borrowers who have recently defaulted turns out to be a non-trivial problem. What are the information structures and strategies that support efficient lending in such an environment?

A natural conjecture is that providing maximal information is best, so that the lender has complete information on the past K outcomes of the borrower, where K is the bound on memory. This turns out to be false. Perfect information on the recent past behavior of the borrower, in conjunction with bounded memory, precludes any lending, because it allows lenders to cherry-pick those borrowers with the strongest long-term incentives to repay. In *any* equilibrium that satisfies a mild and realistic requirement of being robust to small payoff perturbations, borrower exclusion unravels.

Specifically, a borrower whose most recent default is on the verge of disappearing from her record has the same incentives as a borrower with a clean record. Thus, she will repay a loan whenever long-term incentives are such that a borrower with a clean record does so. Lenders, who are able to distinguish her from more recent defaulters, find it profitable to extend her a loan, thereby reducing the length of her punishment. Repeating this argument, by induction, no length of punishment can be sustained. As a result, no lending can be supported. The key problem is that, under perfect information, lenders cannot be disciplined to not make loans to borrowers with a bad record.

This negative result leads us to explore information structures that provide the lender with simple, binary information about the borrower' history. Specifically, the lender is told

[4]Indeed, in a modern economy, borrowers have access to a variety of sources of financing.

[5]The reader may wonder why the lender cannot be disciplined by allowing future borrowers to condition their behavior on the lender's current decision. This mechanism, which is standard in many repeated games, turns out to be unviable in our setting, due to the bounded memory constraint. This is discussed in Section 7.4.

only whether the borrower has ever defaulted in the past K periods (labelled a *bad credit history*) or not (labelled a *good credit history*). A borrower's long-term incentives to repay a new loan differ according to the most recent instance of default in her history. More recent defaulters, with most of their exclusion phase ahead of them, have a stronger incentive to recidivate. Since lenders do not have precise information on the timing of defaults, they are unable to target their loans to defaulters who are more likely to repay.

Coarse information therefore generates *endogenous adverse selection* among the pool of borrowers with a bad credit history, thereby mitigating the tendency of the lender to undermine punishments. To our knowledge, this is the first paper that leverages endogenous adverse selection in order to address an underlying moral hazard problem without adverse selection. Our question is, how can coarse information and the consequent adverse selection be tailored to sustain efficient outcomes?

The simple binary information structure just described prevents a total breakdown of lending. If the punishment phase is sufficiently long, the pool of lenders with a bad credit history is sufficiently likely to re-offend, on average, as to dissuade rogue loans by the lender. Depending on the (exogenous) profitability of loans, the length of exclusion may be longer than is needed to discipline a defaulting borrower. Indeed, disciplining the lender to not lend to borrowers with a bad credit history may require longer punishments than those that suffice to deter a borrower with a good credit history from defaulting.

Nonetheless, we show that under the simple binary information structure, there always exists an equilibrium where the borrower exclusion is minimal, so that borrower payoffs are constrained optimal, subject to integer constraints. If loans are not very profitable, then this achieved in a pure strategy equilibrium and the lender's profits are also constrained optimal. If loans are very profitable, then the equilibrium with minimal exclusion requires that borrowers with bad credit histories be provided loans with positive probability. Some of them will default, altering the constitution of the pool of borrowers with bad credit histories, as borrowers with stronger incentives to re-offend will be over-represented. This serves to discipline lenders. Paradoxically, if individual loans are very profitable, an equilibrium with random exclusion may result in low overall profits for lenders, by inducing a large pool of borrowers with bad records, even though borrower payoffs are high.

Thus, the simple binary information structure can always ensure constrained efficient payoffs for the side of the market that is subject to moral hazard (the seller of the good when quality is variable, or the borrower), but may lead to low payoffs for the side that is not subject to moral hazard (the consumer of the good or the lender).

Next, we show that payoffs on both sides of the market can attain the constrained optimal level under a non-monotonic information partition, where borrowers with multiple defaults are treated favorably and pooled with non-defaulters. Lenders are disciplined because all past defaulters are provided with strong incentives to re-offend.

The non-monotonic information structure might be unappealing in some contexts. We therefore go on to explore other ways in which payoffs can be improved for the side of the market that is not subject to moral hazard (the lender). First, we consider the case where there are multiple Pareto-efficient outcomes in the stage game. In the lender-borrower example, this might correspond to loans different interest rates. We show that supporting an outcome that favours the party subject to moral hazard may increase equilibrium payoffs for both parties. We also examine more complex versions of the stage game. If the party to be trusted (the borrower) has to initiate the interaction, incurring a small cost, this ensures full efficiency. In conjunction with coarse information, such a modification transforms the interaction between borrowers and lenders into a signaling game. Among the borrowers with a bad credit history, those who intend to default have stronger incentives to apply than those who intend to repay. Consequently, lenders are suspicious of applicants with a bad record. Finally, we show that our results apply to any sequential-move stage game where moral hazard is effectively one-sided.

The remainder of this section discusses the related literature. Section 2 sets out the model. Section 3 derives the constrained efficiency benchmarks, which can be attained with infinite memory. It also shows that with bounded perfect memory, no lending can be supported. Section 4 describes our information design problem, and shows that a simple, binary information structure prevents the breakdown of lending. Section 5 shows that such an information structure ensures constrained efficient payoffs for the borrower, either via pure strategies or mixed strategies. Section 6 examines the role of non-monotone information structures in disciplining lenders. Section 7 presents several extensions. The final section concludes.

1.1 Related Literature

Our paper is most closely related to the literature on repeated games with community enforcement. In Kandori (1992), Ellison (1994) and Deb (2008), players belonging to a small (finite) population are randomly matched in each period to play the prisoner's dilemma. A key feature of the analysis is that contagion strategies, where a single defection results in the breakdown of cooperation throughout the population, are used in order to support cooperation. Deb and González-Díaz (2010) extend this analysis to more general simultaneous move stage games.[6]

We assume a large (continuum) population where contagion strategies cannot be effective.[7] Thus our paper is more closely related to Takahashi (2010) and Heller and Mohlin (2017), who analyze the prisoner's dilemma played in a large population, and assume that each player observes an aspect of the previous history of her opponent. Takahashi (2010) shows that if each player observes the entire sequence of past actions taken by her opponent, or observes the *action profile* played in the previous period by her opponent and her opponent's partner, then cooperation can be supported by using "belief-free" type strategies, where a player is always indifferent between cooperating and defecting.[8] He also shows that grim-trigger strategy equilibria sustain cooperation when only the partner's action in the previous period is observed, if the prisoner's dilemma game is supermodular, but not if it is submodular. Heller and Mohlin (2017) assume that a player observes a random sample of the past actions of her opponent. They assume that a small fraction of players are commitment types, an assumption that enables them to rule out belief-free strategies. They show that cooperation can be supported if the prisoner's dilemma payoffs are supermodular but not if they are submodular. Heller and Mohlin (2017) share our concern that equilibria should be robust — while they invoke commitment types to rule out belief-free strategies, we use a purification argument.

Our main departure from the existing literature on community enforcement is that we consider stage games with a sequential structure. This makes a considerable difference to the analysis: moral hazard is one-sided. To illustrate, in our trust game, only the borrower has an incentive to deviate if both players expect the efficient outcome to be played. This feature is not specific to the trust game, but arises from the sequential structure. Indeed, the prisoner's dilemma with sequential moves, or *any game where each player moves at most once*, inherits this feature. The sequential structure also implies that, unlike Takahashi (2010) or Heller and Mohlin (2017), our substantive results do not depend on whether the

[6] Nava and Piccione (2014), Wolitzky (2012) and Ali and Miller (2013) analyze community enforcement where the interaction structure is determined by a network.

[7] Experimental evidence suggests that, for contagion strategies to work, the societies must be very small — Duffy and Ochs (2009) find that cooperation is hard to sustain under random matching, even when the society consists of only 6-10 individuals, while the positive results in Camera and Casari (2009) are for societies consisting of four individuals.

[8] The two cases are closely related to the belief-free strategies considered in Piccione (2002) and Ely and Välimäki (2002) respectively. Observe that finite memory precludes a player observing the entire history of actions taken by his opponent. Finite memory belief-free strategies in our setting are not purifiable.

stage game payoffs are supermodular or submodular.[9]

A second feature of our analysis that departs from most of the literature on community enforcement is that we explicitly assume imperfect monitoring, with some defaults being unavoidable. Thus, equilibria can never be fully efficient, and our focus is on constrained efficiency, where borrower exclusion is temporary.

Finally, we require that equilibria be robust to small payoff shocks, and therefore be purifiable, as in Harsanyi (1973). We view purifiability as a mild robustness requirement. In our context, purifiability rules out equilibria that have the flavour of belief-free equilibria. While belief-free equilibria play a major role in establishing a folk-theorem in repeated games with private monitoring (see Sugaya (2013)), it can be argued that they are unrealistic, and purification arguments are a way of making this argument precise. Our substantive results differ markedly from the negative results in Bhaskar (1998) and Bhaskar, Mailath, and Morris (2013), which demonstrate that purifiability, in conjunction with bounded memory, results in a total breakdown of cooperative behavior. In contrast, the present paper shows that, by providing partial information on past histories, one can robustly support efficient outcomes. It is noteworthy that both the information structures and equilibrium strategies that sustain efficiency are extremely simple and also intuitive.

We assume that the relationship between any pair of individuals is short-lived, so that long-term incentives can only be provided if subsequent partners have information on past behavior. This distinguishes our setting from efficiency wage type models, where the relationship is potentially long-lived, but where a deviating party has the option of starting a new relationship. Dutta (1992) and Kranton (1996) analyze the prisoner's dilemma played in such an environment, and show that new relationships must include an initial non-cooperative phase of "starting small". Ghosh and Ray (1996) point out that the initial phase of starting small is not renegotiation-proof, but that exogenous adverse selection alleviates the problem, by making the initial phase renegotiation-proof. A novel feature of our analysis is the role of *endogenous adverse selection*. Our underlying environment has moral hazard but no adverse selection. We find that an optimal information structure does not fully reveal the borrower's recent history to the lender, thereby endogenously generating adverse selection.

One-sided moral hazard has been studied in the large literature on seller reputation. Most closely related is Liu and Skrzypacz (2014), who assume that buyers are short-lived and have bounded information about the seller's past decisions, but do not observe the magnitude of

[9]With simultaneous moves, each player's best response in an interaction depends upon her expectations regarding the other player's action. With sequential moves, this is not the case for the last mover, and thus inductive reasoning plays a major role.

past sales, and are therefore not able to infer the information observed by past buyers. Buyers also assign a small probability to the seller being committed to high quality. Since the normal type of seller has a greater incentive to cheat when sales are larger, equilibria display a cyclical pattern, whereby the seller builds up his reputation before milking it.[10] Ekmekci (2011) studies the interaction between a long-run player and a sequence of short-run players, where the long-run player's action is imperfectly observed, and there is initial uncertainty about the long-run player (as in reputation models). He shows that bounded memory allows reputations to persist in the long run, even though they necessarily dissipate when memory is unbounded.

Our work also relates to the burgeoning literature on information design, initiated by Kamenica and Gentzkow (2011), and pursued by Kremer, Mansour, and Perry (2014), Bergemann and Morris (2016), and Ely (2017), among others. While this literature has focused on the case of one or few players, our design question relates to a large society. Whereas the distribution of types or states is usually exogenous in the information design context, the induced distributions over types (or private histories) in our paper arise endogenously, as a by-product of the information structure itself.

Our work also relates to the influential macroeconomics literature on "money and memory". Kocherlakota (1998) shows that money and unbounded memory play equivalent roles. Wiseman (2015) demonstrates that, when memory is bounded, money can sustain greater efficiency than memory can.

A key assumption in our analysis is the assumption of bounded memory. While we believe this to be plausible in any context, it is legally enforced in credit markets, where bankruptcy flags have to be removed from borrowers records after a fixed length of time. One might ask whether the existing credit scoring system gets around this legal requirement, and conjecture that FICO scores "remember" dropped bankruptcy flags. However, the empirical evidence presented in Musto (2004), Gross, Notowidigdo, and Wang (2016) and Dobbie, Goldsmith-Pinkham, Mahoney, and Song (2016) shows that this is not the case. A shared finding is that the removal of bankruptcy flags leads to a large jump in credit scores and a large increase in the consumer's access to credit. It is implausible to attribute such a large jump to a sudden, dramatic change in a consumer's defaulting behaviour. Instead, it suggests that bounded memory constraints are binding, and information does indeed disappear when default flags are dropped.[11] Our analysis raises new empirical questions that are discussed

[10] Sperisen (2016) extends this analysis by considering non-stationary equilibria.

[11] For example, Dobbie, Goldsmith-Pinkham, Mahoney, and Song (2016)) find that the increase in credit scores in the quarter of removal of the bankruptcy flag corresponds to an implied 3 percentage point reduced

in our concluding section.

Finally, there is a theoretical literature on credit markets that argue that limited records may be welfare-improving in the presence of adverse selection. This includes Elul and Gottardi (2015) and Kovbasyuk and Spagnolo (2016).

2 The Model

Time is discrete and the horizon infinite. In each period, individuals from a continuum (male) population 1 are randomly matched with individuals from a continuum (female) population 2 to play a sequential-move game Γ.

2.1 The Stage Game

The game Γ is defined as follows.[12] Player 1 chooses an action a_1 from a finite set A_1. Next, nature determines, according to a lottery that may depend on a_1, whether player 2 moves or nature itself moves again. If player 2 is called upon to move, she observes the choice made by player 1, and chooses an action from a finite set $A_2(a_1)$.[13] If nature moves again after a_1, she chooses from the same set $A_2(a_1)$ according to a fixed probability distribution. Let Z denote the terminal nodes of the game Γ. A typical element z consists of a triple (a_1, a_2, i), $i \in \{0, 2\}$, and specifies the actions taken, as well as whether the second action was chosen by player 2 or by nature. For $i \in \{1, 2\}$, $u_i : Z \to \mathbb{R}$ specifies the payoff of player i at each terminal node, and $u_i(y)$ denotes the expected utility of a lottery $y \in \Delta(Z)$.

Fix a pair of actions $a = (a_1, a_2)$ where $a_1 \in A_1$ and $a_2 \in A_2(a_1)$. The *outcome* of a is the distribution over the terminal nodes in Z that results when a_1 is chosen by player 1 and a_2 is chosen by player 2 when she is called upon to move. Since nature could be called upon to move after a_1, and choose an action different from a_2, the outcome could be random. Generalizing this definition, the outcome of a pure strategy profile $\sigma = (a_1, \sigma_2)$, with $\sigma_2 : A_1 \to A_2$ denoting player 2's strategy, is the outcome of $(a_1, \sigma_2(a_1))$.

We make the following assumption on the payoffs of the stage game, Γ, that is satisfied generically:

default risk, on a pre-flag removal risk of 32 percent.

[12] All our results can be understood in the context of the trust game, illustrated in Figure 1, and some readers may wish to skip this section and proceed directly to Section 2.2.

[13] After some choices of player 1, player 2 may not get to move, so that $A_2(a_1)$ is the empty set.

Assumption 1 *There is a unique backwards induction strategy profile,* $\bar{\sigma} = (\bar{a}_1, \bar{\sigma}_2)$*, with outcome* $\bar{y}$*.*

Suppose that there exists a pair of actions $a_1^* \in A_1$ and $a_2^* \in A_2(a_1^*)$ such that the resulting outcome y^* is Pareto-efficient, and also strictly Pareto-dominates $\bar{y}$, i.e. $u_i(y^*) > u_i(\bar{y})$ for every $i \in \{1, 2\}$. (If there is no such pair, then the backwards induction outcome is Pareto-efficient and community enforcement is moot.) Define $\sigma^* = (a_1^*, \sigma_2^*)$ as follows. The strategy σ_2^* plays a_2^* after a_1^*, and the backwards induction strategy $\bar{\sigma}_2$ after any other action by player 1. Clearly, σ^* implements the Pareto-efficient outcome y^*. The following lemma, which is proved in Appendix A.1, shows that, along the path to a Pareto-efficient outcome, in any game where each player moves at most once, moral hazard is one-sided: only the player who moves second has an incentive to deviate.

Lemma 1 *Let* Γ *be a two player game where, along any path of play, each player moves at most once. Let* y^**, the outcome of* σ^**, strictly Pareto-dominate the backwards induction outcome* $\bar{y}$*. Only player 2 has an incentive to deviate from* σ^**, and her optimal deviation is to play* $\bar{\sigma}_2$*.*

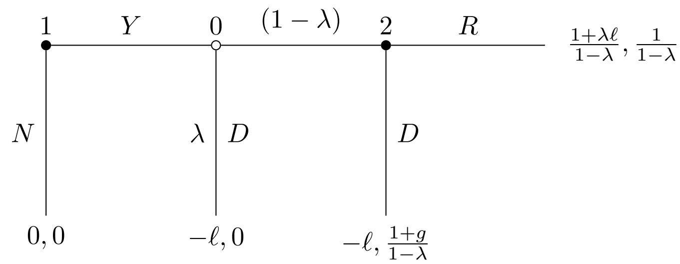

(a) *Extensive Form.*

	R	D
Y	$1, 1$	$-\ell, 1+g$
N	$0, 0$	$0, 0$

(b) *Strategic form.*

Figure 1: *Extensive and strategic form representations of the Trust Game*

Our leading example is the trust game, illustrated in Figure 1a. Player 1 moves first, choosing whether to trust (Y) player 2 or not (N). If he chooses N, the game ends, and both parties get a payoff of zero. If he chooses Y, then player 2 must decide whether to repay this trust (R), or to default (D). However, with a small probability λ, player 2 is unable to repay trust, i.e. she is constrained to default. In this game, it is profitable for player 1 to trust player 2 if the latter intends to repay, and unprofitable if she intends to default. Moreover, wilful default is profitable for player 2. The strategic form of the game, given in Figure 1b, clarifies the players' incentives. Since $g > 0$ and $\ell > 0$, the strategic form is a one-sided prisoner's dilemma, where it is optimal for player 1 to trust if he expects that player 2 will repay the trust when possible, and where player 2 prefers to wilfully default if she is trusted. The key features of the trust game are as follows:

- The outcome of the backwards induction profile $\bar{\sigma} = (N, D)$, where player 1 chooses N and player 2 chooses D, is inefficient, and Pareto-dominated by the (random) outcome that results when the players play $\sigma^* = (Y, R)$.

- If both players expect $\sigma^* = (Y, R)$ to be played, only player 2 has an incentive to deviate.

The trust game has many economic interpretations. In the first, player 1 is the buyer of a product, and 2 is the seller, who must decide whether to supply high quality or low quality, in the event that 1 makes a purchase. However, even if the seller decides to supply high quality, realized quality might turn out to be low. In the second interpretation, player 1 is a lender, and player 2 a borrower. R corresponds to repaying the loan, while D corresponds to defaulting. Lending is profitable if the borrower intends to repay when able; however, there is some probability that the borrower is not able to repay even if she wants to.

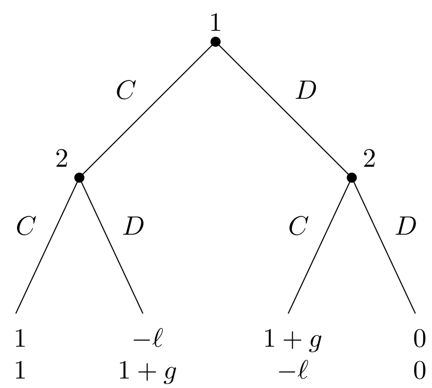

(a) *Extensive Form.*

	σ_2^*	$\bar{\sigma}_2$
σ_1^*	$1, 1$	$-\ell, 1+g$
$\bar{\sigma}_1$	$0, 0$	$0, 0$

(b) *Reduced strategic form.*

Figure 2: *Extensive and reduced strategic form representations of the Prisoner's Dilemma*

Our second example is the sequential move prisoner's dilemma, depicted in Figure 2a. The backwards induction strategy profile has both players always choosing D, resulting in the payoffs $(0, 0)$. Consider the strategy profile σ^* where player 1 plays C and player 2 responds to C with C and to D with D; this results in the payoff vector $(1, 1)$. A reduced strategic form of the game highlighting these strategies is depicted in Figure 2b.[14] For player

[14] We use the term *reduced strategic form* to denote a specific 2×2 strategic-form game emphasising the strategies of interest — this differs from the way in which the term is often used in the literature.

1, $\sigma_1^* = C$ and $\bar{\sigma}_1 = D$. For player 2, the strategy σ_2^* represents C after C and D after D, while the backwards induction strategy $\bar{\sigma}_2$ always plays D. Observe that player 2 has an incentive to deviate from σ^*, but player 1 does not — if he plays D, then player 2 responds with D, and his payoff is 0. Imperfect monitoring of player 2's actions can be introduced by assuming that, with probability λ, nature chooses player 2's action to be D. As we will see, the formal analysis is then identical to that of the trust game.

We henceforth restrict attention to games Γ where each player moves once, and where the backwards induction outcome is Pareto-dominated by the outcome of (a_1^*, a_2^*). For the sake of exposition, we make the following payoff normalizations. Let the payoffs from the backwards induction outcome be $(0, 0)$, the payoffs from the outcome of (a_1^*, a_2^*) be $(1, 1)$, and the payoffs from the outcome of $(a_1^*, \bar{\sigma}_2(a_1^*))$ be $(-\ell, 1+g)$. Then, in the class of games we consider, only two strategies will be of interest for each player, σ_i^* and $\bar{\sigma}_i$. The associated reduced strategic form is given in Figure 2b.

Note that the Pareto-efficient profile that Pareto-dominates the backwards induction profile need not be unique, even with generic payoffs.[15] If so, the players will have opposed preferences over these Pareto-efficient profiles. Our analysis speaks to the sustainability of any Pareto-efficient profile. Games with multiple efficient profiles are discussed in Section 7.1.

We extend the analysis beyond games where each player moves at most once. In Section 7.2, we show that if the player who is subject to moral hazard has to initiate the interaction, then sustaining efficient outcomes becomes easier. In Section 7.3 we show that our analysis extends to more general games, where, along the path to the efficient outcome, only one player has the incentive to deviate, given that players respond to any deviation with the backwards induction strategies.

2.2 Information on Stage-Game Outcomes

Our analysis incorporates imperfect monitoring of the actions of player 2 in the repeated game, which is why we allow for a move by nature in the stage game, after player 1 moves. As we will see, imperfect monitoring of player 2's actions imposes a direct incentive-cost — player 2 will need to be punished for transgressions that she did not intend, and this will bound her payoffs away from full efficiency. Imperfect monitoring of player 1's actions

[15]Nonetheless, only player 2 has an incentive to deviate from the path to any Pareto-efficient outcome. The normalization of payoffs in the previous paragraph depends upon the outcome that we are seeking to support.

does not have this drawback, since player 1 has no incentive to deviate on the path to the efficient outcome, which is why we do not consider the possibility that nature chooses player 1's actions.

Let $\mathcal{P}$ be a partition of Z, the set of terminal nodes of Γ, with the following interpretation. Suppose that terminal node z is realized, and that it belongs to the set denoted $\mathcal{P}(z)$. Then any outside observer can (at most) know that an element of $\mathcal{P}(z)$ was realized.

Assumption 2 *For every profile* (a_1, a_2), *the terminal nodes* $z = (a_1, a_2, 0)$ *and* $z' = (a_1, a_2, 2)$ *are such that* $\mathcal{P}(z) = \mathcal{P}(z')$.

That is, no outside observer can distinguish whether the action a_2 was taken by nature or by player 2. In the trust game, an outside observer cannot distinguish involuntary defections from voluntary ones. Thus, the terminal nodes $(Y, D, 0)$ and $(Y, D, 2)$ belong to the same element of $\mathcal{P}$, which we denote $\mathcal{D}$. Letting $\mathcal{N} := \{(N)\}$ and $\mathcal{R} := \{(Y, R, 2)\}$, the information partition of the outside observer in the trust game is $\mathcal{O} := \mathcal{P} = \{\mathcal{N}, \mathcal{D}, \mathcal{R}\}$.

In the general game, Γ, consider the following partition of Z, which we also denote by $\mathcal{O}$. Let $\mathcal{R}$ denote the set of terminal nodes where a_1^* and a_2^* are played, $\mathcal{D}$ denote the set of terminal nodes where a_1^* and $a_2 \neq a_2^*$ are played, and $\mathcal{N}$ denote the set of terminal nodes where $a_1 \neq a_1^*$ is played. Observe that this partition satisfies Assumption 2, but is possibly coarser than entailed by the assumption. For example, in the prisoner's dilemma in Figure 2a, the nodes (D, C) and (D, D) are pooled in the element $\mathcal{N}$, while $\mathcal{R} = \{(C, C)\}$ and $\mathcal{D} = \{(C, D)\}$. In this paper, we shall focus on equilibria where strategies are measurable with respect to the partition $\mathcal{O}$ of the stage game outcomes. This is for expositional convenience, as it enables us to discuss in a unified manner all stage games taking the general form Γ.

Let Γ^∞ denote the infinitely repeated game where at every period players are randomly matched to play the stage game Γ.[16] We assume that player 2 has a discount factor $\delta \in (0, 1)$. The discount factor of player 1 is irrelevant for positive analysis.[17] Since player 2 has a short-term incentive to default, incentives to repay can only be provided by her future partners. Those partners will have some information about her past behavior. The precise detail of their information structure will vary in the next sections, depending on which specification we study.

We focus on stationary Perfect Bayesian Equilibria, where agents are sequentially rational at each information set, and their beliefs are given by Bayes' rule wherever possible. The

[16] The game Γ^∞ will depend on an information structure which is as yet unspecified.

[17] As we will see in Section 7.4, incentives for player 1 have to be provided within the period.

stationarity assumption implies that players do not condition on calendar time. We shall focus on equilibria where all players in population 1 follow the same strategy, and all players in population 2 follow the same strategy. We also require that our equilibria be purifiable, as we now explain.

2.3 Payoff Shocks: The Perturbed Game

We now define $\Gamma(\varepsilon)$, a perturbed version of the extensive form stage game, Γ, indexed by ε, a scaling parameter. Let X denote the set of player decision nodes in Γ and let $\iota(x)$ denote the player who moves at $x \in X$, making a choice from a non-singleton set, $A(x)$. At each such decision node $x \in X$, player $\iota(x)$'s payoff from action $a^k \in A(x)$ is augmented by εz_x^k, where $\varepsilon > 0$. The scalar z_x^k is the k^{th} component of z_x, where $z_x \in \mathbb{R}^{|A(x)|-1}$ is the realization of a random variable with bounded support. We assume that the random variables $\{Z_x\}_{x\in X}$ are independently distributed, and that their distributions are atomless. Player $\iota(x)$ observes the realization z_x of the shock before being called upon to move. In the repeated version of the perturbed game, $\Gamma^\infty(\varepsilon)$, we assume that the shocks for any player are independently distributed across periods.[18] In the buyer-seller interpretation of the trust game, we may assume that the buyer gets an idiosyncratic payoff shock from his outside option of not buying, while the seller gets an idiosyncratic shock to her cost of supplying high quality. Motivated by Harsanyi (1973), we focus on purifiable equilibria, i.e. equilibria of the game without shocks, Γ^∞, that are limits of equilibria of the game $\Gamma^\infty(\varepsilon)$ as $\varepsilon \to 0$.

Harsanyi's purification argument is widely regarded as the most compelling justification for mixed strategy equilibria, since a game with no payoff shocks is an idealization. Harsanyi focused on strategic form games, and showed that for generic payoffs, all Nash equilibria are purifiable, and can be approximated by strict equilibria of the perturbed game. Bhaskar, Mailath, and Morris (2013) show that this is not the case in repeated or stochastic games — in a large class of stochastic games, purifiability refines the set of equilibria, by eliminating belief-free type equilibria.

Call an equilibrium of the unperturbed game *sequentially strict* if a player has strict incentives to play her equilibrium action at every information set, whether this information set arises on or off the equilibrium path. The following lemma, proved in Appendix A.6.3, will prove useful in our subsequent analysis.

Lemma 2 *Every sequentially strict equilibrium of* Γ^∞ *is purifiable.*

[18]The assumption that the lender's shocks are independently distributed across periods is not essential.

2.4 The Trust Game

We now argue that the analysis of the game Γ^{∞}, when the stage game Γ belongs to the class described in Section 2.1 and the strategies of population 1 are measurable with respect to the partition $\mathcal{O}$ of stage-game outcomes, is equivalent to the analysis of the game Γ^{∞}, when the stage game is the trust game given in Figure 1a, with the information partition $\mathcal{O}$.

Lemma 3 *Consider an interaction at date t, between two players from populations 1 and 2, labelled $\hat{1}$ and $\hat{2}$ respectively. Suppose that the strategies in Γ^{∞} of all other players in population 1 are measurable with respect to $\mathcal{O}$. After any history, h^t, player $\hat{1}$'s optimal strategy in the stage game at date t belongs to $\{\bar{\sigma}_1, \sigma_1^*\}$ and player $\hat{2}$'s optimal strategy at date t belongs to $\{\bar{\sigma}_2, \sigma_2^*\}$.*

Lemma 3 shows that, in the game Γ^{∞}, when other player's strategies are measurable with respect to the partition $\mathcal{O}$, only two strategies of the stage game Γ will be of interest for each player: σ_i^* and $\bar{\sigma}_i$. The associated reduced strategic form is given in Figure 2b. Observe that this is the (non-reduced) strategic form associated with only one extensive form game: the trust game.[19] Therefore, it will be sufficient to limit our analysis to the trust game. In keeping with the credit market interpretation, we will refer to players in population 1 as lenders, and those in population 2 as borrowers.

3 Benchmarks

3.1 The Infinite Memory Benchmark

Suppose that each lender can observe the entire history of transactions of each borrower he is matched with. That is, a lender matched with a borrower at date t observes the outcomes in $\mathcal{O}$ of the borrower in periods $1, 2, .., t-1$. We assume that payoff parameters are such that there exists an equilibrium where lending takes place.[20] Assume also that the borrower does not observe any information about the lender, so that incentives for the lender have to be provided within the period.

[19] In the trust game, whose extensive form is illustrated in Figure 1a, player 1 has only two strategies, $\sigma_1^* = Y$ and $\bar{\sigma}_1 = N$, and player 2 has only two strategies, $\sigma_2^* = R$ and $\bar{\sigma}_2 = D$.

[20] That is, we assume that permanent exclusion is sufficiently costly that the Bulow and Rogoff (1989) problem, whereby a lender always finds it better to default and re-invest the sum, does not arise. For example, costs of filing for bankruptcy could be non-trivial. The precise condition is $g < \frac{\delta(1-\lambda)}{1-\delta(1-\lambda)}$.

Consider an equilibrium where a borrower who is in good standing has an incentive to repay when she is able to. Her expected gain from intentional default is $(1-\delta)g$.[21] The deviation makes a difference to her continuation value only when she is able to repay, i.e. with probability $1-\lambda$. Suppose that after a default, wilful or involuntary, she is excluded from the lending market for K periods. The incentive constraint ensuring that she prefers repaying when able is then

$$(1-\delta)g \leq \delta(1-\lambda)[V^K(0) - V^K(K)], \tag{1}$$

where $V^K(0)$ denotes her payoff when she is in good standing, and $V^K(K)$ her payoff at the beginning of the K periods of punishment. These are given by

$$V^K(0) = \frac{1-\delta}{1-\delta[\lambda\delta^K + 1 - \lambda]}, \tag{2}$$

$$V^K(K) = \delta^K V^K(0). \tag{3}$$

The most efficient equilibrium in this class has K large enough to provide the borrower incentives to repay when she is in good standing, but no larger. Call this value $\bar{K}$, and assume that the incentive constraint (1) holds as a strict inequality when $K = \bar{K}$ — this assumption will be made throughout the paper, and is satisfied for generic values of the parameters (δ, g, λ). The payoff of the borrower when she is in good standing is $\bar{V} := V^{\bar{K}}(0)$, i.e. it is given by equation (2) with $K = \bar{K}$. We evaluate the payoffs of any lender by his per-period payoff in the steady state corresponding to this equilibrium. Since the lender earns an expected payoff of 1 on meeting a borrower in good standing, and 0 otherwise, his payoff $\bar{W}$ equals the fraction of borrowers in good standing, i.e. $\bar{W} = \frac{1}{1+\lambda\bar{K}}$.

It is useful at this point to examine the incentives of the lender, given that future play cannot be conditioned on his behavior. We want to ensure that a borrower who defaults, and who should be excluded for K periods, is not offered a loan. To do this, we must distinguish between defaults that occur when a loan should be made, and those that arise when the lender should not have lent in the first place. This is illustrated in the equilibrium described by the automaton in Figure 3, where a defaulting borrower is excluded for K periods — the figure depicts the case of $K = 2$. Depending on the entire history, the borrower is either in a good state or in one of K distinct bad states. The lender extends a loan if and only if the borrower is in the good state. A borrower begins in the good state, and stays there

[21] Per-period payoffs are normalized by multiplying by $(1-\delta)$.

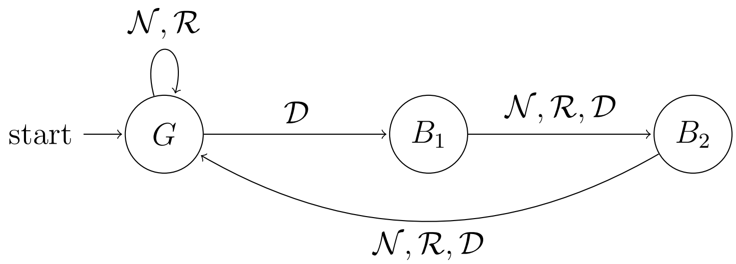

Figure 3: *Strategy profile with two periods of exclusion.*

unless she defaults, in which case she transits to the first of the bad states. The borrower then transits through the remaining $K - 1$ bad states, spending exactly one period in each, and then back to the good state. The transition out of any bad state is *independent of the outcome in that period*, thus ensuring that the borrower's actions in a bad state do not affect her continuation value. Since the borrower is never punished for a default when she is in a bad state, she will always choose to default, ensuring that no lender will lend to her when she is in a bad state. Note the importance of infinite memory: this equilibrium requires that the lender should be able to observe the entire history of outcomes in $\mathcal{O}$ of every borrower he is matched with. Otherwise, he cannot deduce whether the borrower defaulted in a period where she was supposed to be lent to, or one in which she was supposed to be excluded.

The equilibrium with $\bar{K}$ periods of exclusion can be improved upon — due to integer constraints, the punishment is strictly greater than what is required to ensure borrower repayment. In Appendix A.3 we show that the highest payoff the borrower can achieve in *any equilibrium* is V^*, given by the following expression[22], and that it is strictly less than 1 due to imperfect monitoring, since voluntary and involuntary defaults cannot be distinguished:

$$V^* = 1 - \frac{\lambda}{1 - \lambda} g. \tag{4}$$

To sustain the equilibrium payoff V^*, we assume that players observe the realization of a public randomization device at the beginning of each period, and that past realizations of the randomization device are also a part of the public history. The payoff V^* can be achieved by the borrower being excluded for $\bar{K} - 1$ periods with probability x^* and for $\bar{K}$

[22]In deriving this bound, we assume that borrower mixed strategies are not observable. If mixed strategies are observable we can sustain a borrower payoff higher than V^*, as in Fudenberg, Kreps, and Maskin (1990). The borrower in good standing must have access to a private randomization device that allows her to wilfully default with some probability, and such defaults are not punished. Furthermore, past realizations of the randomization device must also be a part of the infinite public history. The assumption that mixed strategies are observable seems strong and possibly unrealistic.

periods with probability $1 - x^*$. This gives rise to a steady-state proportion of borrowers in good standing equal to $\frac{1}{1+\lambda(\bar{K}-x^*)}$, and since the lender gets a payoff of 1 whenever he meets a borrower in good standing, and 0 otherwise, this proportion equals the lender's expected payoff, W^*.

Observe that these payoff bounds generalize beyond the trust game. Consider a general game Γ, and an efficient outcome corresponding to the action profile (a_1^*, a_2^*). Suppose that after a_1^*, with probability λ nature moves, choosing the action $\bar{\sigma}_2(a_1^*)$. With probability $1-\lambda$, player 2 moves. Then the efficiency bounds for player 2 in Γ^∞ are exactly as in the trust game, i.e. $\bar{V}$ and V^*.

To summarize: $\bar{V}$ and $\bar{W}$ will be called the constrained efficient payoffs for the borrower and lender respectively, that reflect both the integer constraint and the incentive constraint under imperfect monitoring. V^* and W^* will be called the fully efficient payoffs — these include the incentive constraint for the borrower, but no integer constraints. We assume that the designer's objective is to achieve a payoff no less than $\bar{V}$ for the borrower. In Section 5 we show that this is always possible, though it sometimes results in low payoffs for the lender. In Sections 6 and 7.2 we show how the designer can correct this, and also achieve $\bar{W}$ for the lender.

3.2 Perfect Bounded Memory

Henceforth, we shall assume that lenders have bounded memory, i.e. we assume that at every stage, the lender observes a bounded history of length K of past outcomes in $\mathcal{O}$ of the borrower he is matched with in that stage. We assume that the lender does not observe any information regarding other lenders. In particular, he does not observe any information regarding the lenders with whom the borrower he currently faces has been matched in the past.

Our first proposition is a negative one — if we provide the lender full information regarding the past K interactions of the borrower, then no lending can be supported.

Proposition 1 *Suppose that $K \geq 2$ is arbitrary and the lender observes the finest possible partition of $\mathcal{O}^K$, or that $K = 1$ and the information partition is arbitrary. The unique purifiable equilibrium corresponds to the lender never lending and the borrower never repaying.*

The proof does not follow directly from Bhaskar, Mailath, and Morris (2013), but is an adaptation of that argument, so we do not present it here. The intuition is as follows. Suppose that the information partition is the finest possible. Consider a candidate equilibrium where a borrower who defaults is excluded for $K \geq \bar{K}$ periods, so that a borrower with a clean record prefers to repay. Consider a borrower with exactly one default which occurred exactly K periods ago. Such a borrower has incentives identical to those of a borrower with a clean record, and will therefore also repay. Therefore, a lender has every incentive to lend to such a borrower, undermining her punishment. An induction argument then implies that no length of punishment can be sustained.

The role of purification is to extend this argument to all possible equilibria. When memory length is K, the borrower knows that the lender tomorrow cannot condition his behavior on events that happened K period ago, since he will not observe these events. The payoff shocks faced by the borrower imply that for any strategies of the lenders, the borrower is indifferent between R and D only on a set of measure zero. Thus, the borrower will also not condition her behavior on what happened K periods ago. In consequence, the lender today will not condition his lending decision on events K periods ago, and an induction argument ensures that there can be no conditioning on history.

Observe that there exist mixed equilibria that support the efficient payoff V^* for the borrower — but these are not purifiable. An example where the lender's strategy conditions only on the borrower outcome in the previous period is as follows. The lender lends with probability one if the borrower's previous outcome is $\mathcal{R}$, i.e. she has repaid, and lends with probability $p < 1$ if the previous outcome is $\mathcal{D}$ or $\mathcal{N}$, where p is chosen so as to make the borrower indifferent between repaying and defaulting. The borrower defaults with probability q every time she has a loan, independent of her previous history, where q makes the lender indifferent between lending and not lending, so that the lender's equilibrium payoff is 0. The lender has identical beliefs about the default probability of the borrower regardless of the lender's outcome in the previous period. However, she lends with certainty after $\mathcal{R}$, but only with probability p otherwise. Consequently, if there are small idiosyncratic shocks that affect the opportunity cost of funds for the lender, then he will condition his behavior on these shocks and not on the previous history of the borrower. Thus, the equilibrium is not purifiable.

The above example clarifies why we insist on purifiability — on our view, equilibria that fail this test are not robust. In this paper, the equilibria we construct will be, for the most part, sequentially strict — every player will always have a strict best response at every information set. We will, on one occasion, consider mixed strategy equilibria, and when we do this, we will prove that these equilibria are indeed purifiable — in the game with payoff shocks, the approximation to the mixed equilibrium we consider will be sequentially strict.

The second part of the proposition, that there cannot be conditioning upon history if $K = 1$, does not need an induction argument and applies to any information structure. So we need $K \geq 2$, since otherwise no information structure sustains lending. Even if $g < \frac{\delta(1-\lambda)}{1+\delta\lambda}$, so that only one period of memory is required to satisfy (1), we need at least two period memory. Underlying this second result is a more general point, that will recur frequently in our analysis — the borrower will never condition her behavior on what happened K periods ago, under *any* information structure.

Total breakdown of lending does not occur under all information structures. Surprisingly, *less information* may support cooperative outcomes, as we shall see below. Although it still remains the case that the borrower will not condition on events that happened exactly K periods ago, a coarser information structure prevents the lender from knowing this, and thus the induction argument underlying the above proposition (and the main theorem in Bhaskar, Mailath, and Morris (2013)) does not apply. Such an information structure generates *endogenous adverse selection*, and the main contribution of this paper is to show this can be used to sustain efficient or near-efficient outcomes, in repeated game type environments where fine information leads to a breakdown of cooperation.

4 Information

We have in mind a designer or social planner who, subject to memory being bounded, designs an information structure for this large society, and recommends a non-cooperative equilibrium to the players.[23] The designer's goal is to achieve a borrower payoff no lower than $\bar{V}$, and a lender payoff no lower than $\bar{W}$ — our focus is mainly on the former. This requires supporting equilibria where lending is sustained, and where a borrower who defaults is not excluded for longer than necessary.

Let K denote the bound on memory chosen by the designer — we allow K to be arbitrarily large but bounded by $\tilde{K}$, and exogenous, finite bound on memory. An information system provides information to the lender based on the past K outcomes in $\mathcal{O}$ of the borrower, with the set of K-period histories being $\mathcal{O}^K$. We assume that the borrower does not receive information on the past outcomes of the lender — in Section 7.4 we show that such information would be useless in any equilibrium, since no borrower would condition on it. Information structures fall into two broad categories.

[23]Thus, the designer cannot dictate the actions to be taken by any agent, and in particular cannot direct lenders to refrain from lending to defaulters.

A *deterministic* information (or signal) structure consists of a finite signal space S and a mapping $\tau : \mathcal{O}^K \to S$. More simply, it consists of a partition of the set of K-period histories, $\mathcal{O}^K$, with each element of the partition being associated with a distinct signal in S, and can also be called a *partitional* information structure.

A *random* information (or signal) structure allows the range of the mapping to be the set of probability distributions over signals, so that $\tau : \mathcal{O}^K \to \Delta(S)$.

The formal definitions are useful in ensuring that one respects the bounded memory constraints. Note that in both cases, the signal *does not* depend on past signal realizations, since otherwise one could smuggle in infinite memory on outcomes. Most of our analysis will focus on partitional information structures, for two reasons. First, we will see that the efficiency gains from random information are restricted to overcoming integer problems, and are therefore of less interest. Second, we deem a random information structure to be more vulnerable to manipulation, especially in a large society — if a borrower with a given history gets a higher continuation value from one signal realization than another, she may wish to expend resources to influence the outcome.

4.1 A Simple Binary Information Structure

We shall assume henceforth that the exogenous bound on the length of memory, $\tilde{K}$, is finite but large enough that it is never a binding constraint. So the length of effective memory, K, can be chosen without constraints.[24] Assume that $K \geq \max\{\bar{K}, 2\}$,[25] and let the information structure be given by the following binary partition of $\mathcal{O}^K$. The lender observes a "bad credit history" signal, B, if and only if the borrower has had an outcome of $\mathcal{D}$ in the last K periods, and observes a "good credit history" signal, G, otherwise. Since this information structure will recur through this paper, it will be convenient to label it the *simple* information partition/structure. In this section we show that lending can be supported under the simple information structure.

The borrower has complete knowledge of her own private history, since she knows the entire history of past transactions. Information on events that occurred more than K periods ago is irrelevant, since no lender can condition on it. Under the simple information structure, the following partition of K-period private histories will be used to describe the borrower's

[24] Observe that K can be set less than $\tilde{K}$ by not disclosing any information about events that occurred more than K periods ago. More subtly, this can also be achieved by *full disclosure* of events that occurred between K and $\tilde{K}$ periods ago — this follows from arguments similar to those underlying Proposition 1.

[25] We also examined how lending can be sustained when $\tilde{K} < \bar{K}$, but for reasons of space do not present these results here.

incentives. Partition the set of private histories into $K+1$ equivalence classes, indexed by m. More precisely, let t' denote the date of the most recent incidence of $\mathcal{D}$ in the borrower's history, and let $j = t - t'$, where t denotes the current period. Define $m := \min\{K+1-j, 0\}$. Under the simple information structure, if $m = 0$ the lender observes G while if $m \geq 1$ the lender observes B. Thus, m represents the number of periods that must elapse without default before the borrower gets a good history. When $m \geq 1$, this value is the borrower's private information. In particular, among lenders with a bad credit history, the lender is not able to distinguish those with a lower m from those with a higher m.

Consider a candidate equilibrium where the lender lends after G but not after B, and the borrower always repays when the lender observes G. Let $V^K(m)$ denote the value of a borrower at the beginning of the period, as a function of m. When her credit history is good, the borrower's value is given by $V^K(0)$ defined in (2). For $m \geq 1$, the borrower is excluded for m periods before getting a clean history, so that

$$V^K(m) = \delta^m V^K(0), \qquad m \in \{1, \ldots, K\}. \tag{5}$$

Since $K \geq \bar{K}$, the borrower strictly prefers to repay at a good credit history. Let us examine the borrower's repayment incentives when the lender sees a bad credit history. Note that this is an unreached information set at the candidate strategy profile, since the lender is making a loan when he should not. Repayment incentives depend upon the borrower's private information, and are summarized by m. Observe that the borrower's incentives at $m = 1$ are identical to those at $m = 0$ — for both types of borrower, their current action has identical effects on their future signal. Therefore, a borrower of type $m = 1$ will always choose R. Now consider the incentives of a borrower of type $m = K$. We need this borrower to default, since otherwise every type of borrower would repay and lending after observing signal B would be optimal for the lender. Thus we require

$$(1-\delta)g > \delta(1-\lambda)\left[V^K(K-1) - V^K(K)\right]. \tag{6}$$

The left-hand side above is the one-period gain from default, whereas the right-hand side reflects the gain in continuation value from repayment, since the length of exclusion is reduced by one period. In Appendix A.4 we show that (6) is satisfied whenever $K \geq \bar{K}$.

Now consider the incentives to repay for a borrower with an arbitrary $m > 1$. By repaying, the length of exclusion is reduced to $m-1$, while by defaulting, it increases to K.

Thus the difference in overall payoffs from defaulting as compared to repaying equals

$$(1-\delta)g - \delta(1-\lambda)[V^K(m-1) - V^K(K)] = (1-\delta)g - (1-\lambda)(\delta^m - \delta^{K+1})V^K(0). \quad (7)$$

We have seen that when $m = K$, the above expression is positive, while when $m = 1$, the expression is negative. Thus there exists a real number, denoted $m^\dagger(K) \in (1, K)$, that sets the payoff difference equal to zero. We assume that $m^\dagger(K)$ is not an integer, as will be the case for generic payoffs.[26] Let $m^*(K) = \lfloor m^\dagger(K) \rfloor$, i.e. m^* denotes the integer value of $m^\dagger$. If $m > m^*(K)$, the borrower chooses D when offered a loan. If $m \leq m^*(K)$, she chooses R. Intuitively, a borrower who is close to getting a clean history will not default, just as a convict nearing the end of his sentence has incentives to behave.

Since the lender has imperfect information regarding the borrower's K-period history, we will have to compute the lender's beliefs about those histories. These beliefs will be determined by Bayes rule, from the equilibrium strategy profile. We will focus on lender beliefs in the steady state, i.e. under the invariant distribution over a borrower's private histories induced by the strategy profile. In Appendix A.5 we set out the conditions under which the strategies set out here are optimal in the initial periods of the game, when the distribution over borrower types may be different from the steady state one.

We now describe the beliefs of the lender when he observes B. In every period, the probability of involuntary default is constant, and equals λ. Furthermore, under the candidate strategy profile, a borrower with a bad credit history never gets a loan and hence transits deterministically through the states $m = K, K-1, .., 1$. Therefore, the invariant (steady state) distribution over values of m induced by this strategy profile gives equal probability to each of these states.[27] Consequently, the lender attributes probability $\frac{m^*(K)}{K}$ to a borrower with signal B repaying a loan. Simple algebra shows that making a loan to a borrower with a bad credit history is strictly unprofitable for the lender if

$$\frac{m^*(K)}{K} < \frac{\ell}{1+\ell}. \quad (8)$$

Suppose that ℓ is large enough that the lender's incentive constraint (8) is satisfied. Then he finds it strictly optimal not to lend after B, and to lend after G. We have seen that, since $K \geq \bar{K}$, it is optimal for the borrower to repay when she has a clean history G. Moreover, if granted a loan after B, she has strict incentives to repay as long as $m \leq m^*(K)$ and to

[26] This ensures that the equilibrium is sequentially strict, permitting a simple proof of purifiability.

[27] The invariant distribution $(\mu_m)_{m=0}^K$ has $\mu_0 = \frac{1}{1+K\lambda}$ and $\mu_1 = \cdots = \mu_K = \frac{\lambda}{1+K\lambda}$.

default if $m > m^*(K)$. Thus, there exists an equilibrium that is sequentially strict, and is therefore purifiable. In other words, providing the borrower with coarse information, so that he does not observe the exact timing of the most recent default, overcomes the impossibility result in Proposition 1. Even though those types of borrowers who are close to "getting out of jail" would choose to repay a loan, the lender is unable to distinguish them from those whose sentence is far from complete. He therefore cannot target loans to the former. In other words, coarse information endogenously generates borrower adverse selection that mitigates *lender moral hazard*.

It remains to investigate the conditions on the parameters that ensure that the incentive constraint (8) is satisfied. In Appendix A.4, we show that $\frac{m^*(K)}{K} \to 0$ as $K \to \infty$. We therefore have the following proposition:

Proposition 2 *An equilibrium where the lender lends after observing G and does not lend after observing B exists as long as K is sufficiently large. Such an equilibrium is sequentially strict and therefore purifiable.*

4.2 Discussion

The previous result highlights the novel role of *endogenous adverse selection* in sustaining efficient outcomes. The simple information structure provides the lender with coarse information about the borrower's outcomes, thereby generating uncertainty about the lender's private history. This prevents the lender from cherry-picking among the borrowers with a bad credit history. Thus there is no unravelling due to a induction argument, as in Proposition 1. This is despite the fact that our underlying economic environment has moral hazard but no adverse selection. Coarse information therefore makes a qualitative difference and plays a more important role as compared with Kamenica and Gentzkow (2011) and the subsequent literature on information design. There, it serves to increases the probability with which the agent takes the action desired by the principal, by pooling states where the agent's incentive to take this action is strict with states where the incentive constraint is violated. In our context, there may be but a single history where the lender has an incentive to supply a loan when she should not. With perfect information, this violation causes unravelling so that no lending can be supported at all. Coarse information, by preventing the lender from detecting this single history, prevents this unravelling.

We note briefly that moderate exogenous adverse selection among borrowers may actually improve matters, by mitigating a lender's incentive to lend at a bad history. In other words, exogenous borrower adverse selection can augment the adverse selection that is endogenously generated by our information structure. Suppose that there are two types of borrowers, who differ only in their rates of involuntary default, λ and λ'. Let $\lambda' > \lambda$, so that the former corresponds a high-risk borrower. Given the payoffs in Figure 1a, the expected payoff of lending to a high-risk borrower who intends to repay is $\pi := \frac{1+(\lambda'-\lambda)\ell-\lambda'}{(1-\lambda)} < 1$. The payoff from lending to a borrower who intends to default is $-\ell$, independent of her type. Let θ denote the fraction of high-risk borrowers. Assume K-period memory, and the simple information structure. Consider a pure strategy profile where lenders lend after signal G, but not after B. Suppose that K is large enough that borrowers of either type find it optimal to repay at credit history G. The steady state probability that a high-risk borrower has a bad credit history equals $\frac{K\lambda'}{1+K\lambda'}$, which exceeds the steady state probability that a normal borrower has signal B, $\frac{K\lambda}{1+K\lambda}$. Thus the probability assigned by the lender to a borrower with a bad signal being high-risk is greater than θ. Since high-risk borrowers are less profitable even when they intend to repay (i.e. when $m \leq m^*$), this reduces the lender's incentive to lend after a bad signal. Thus exogenous borrower adverse selection mitigates the lenders' tendency for rogue lending, even though we have shown that it is not essential.

5 Efficiency Under the Simple Information Structure

In this section we investigate the conditions under which an equilibrium with punishments of minimal length, $\bar{K}$, exists, for $\bar{K} \geq 2$, under the simple information structure. We show that for all parameter values, the borrower's constrained efficient payoff $\bar{V}$ can always be achieved.

5.1 Pure strategy equilibrium when ℓ is large

Consider first the pure strategy profile set out in the previous section, with $\bar{K}$ memory. In the appendix, in lemma A.1 we show that $m^*(\bar{K}) = 1$, so that every borrower with type $m > 1$ defaults, giving rise to a steady state repayment probability of $\frac{1}{\bar{K}}$. The lender has strict incentives not to lend to a borrower with a bad credit history if

$$\frac{1}{\bar{K}} < \frac{\ell}{1+\ell} \Leftrightarrow \ell > \frac{1}{\bar{K}-1}. \tag{9}$$

Given that punishments are of length $\bar{K}$, a borrower with a good credit history has a strict incentive to repay. Thus we have a sequentially strict equilibrium that achieves the payoff

$\bar{V}$ for the borrower and $\bar{W}$ for the lender. We can also achieve the fully efficient payoffs V^* and W^* by using a random signal structure, to induce a punishment length between $\bar{K}$ and $\bar{K}-1$. We define the *random version of the simple information structure* as follows. If there is no instance of $\mathcal{D}$ in the last $\bar{K}$ periods, signal G is observed by the lender. If there is any instance of $\mathcal{D}$ in the last $\bar{K}-1$ periods, then signal B is observed. Finally, if there is a single instance of $\mathcal{D}$ in the last $\bar{K}$ periods and this occurred exactly $\bar{K}$ periods ago, signal G is observed with probability $(1-x)$, and B is observed with probability x. We assume $x > x^*$, where x^* denotes the value where the borrower is indifferent between repaying and defaulting when she has signal G. In Appendix A.4 we show that under this random signal structure, $m^* = 1$, so that the repayment probability after a bad signal remains low enough and lending is not profitable, thereby proving the following proposition.

Proposition 3 *Suppose $\bar{K} \geq 2$. If loans are not too profitable, so that $\frac{1}{\bar{K}} < \frac{\ell}{1+\ell}$, there exist sequentially strict equilibria that can a) achieve constrained efficient payoffs $\bar{V}$ and $\bar{W}$ under the simple information structure, and b) approximate the fully efficient payoffs V^* and W^* under a random signal structure.*

5.2 Mixed equilibrium when ℓ is small

Consider now the case where $\ell < \frac{1}{\bar{K}-1}$. Suppose that lenders lend with positive probability on observing B, and lend with probability one after G.[28] We now show that this permits an equilibrium where the length of exclusion after a default is no greater than $\bar{K}$ — indeed, the effective length is strictly less, since exclusion is probabilistic. This may appear surprising — if a lender is required to randomize after B, then not lending must be optimal, and so the necessary incentive constraint for an individual lender should be no different from the pure strategy case. However, the behavior of the population of lenders changes the relative proportions of different types of borrower among those with signal B. It raises the proportion of those with larger values of m, thereby raising the default probability at B and disciplining lenders. Thus, other lenders lending probabilistically to borrowers with a bad history exacerbates the adverse selection faced by the individual lender.

Let $p \in (0, 1)$ denote the probability that a borrower with history B gets a loan (a borrower with history G gets a loan for sure). Recall that if $p = 0$ and $K = \bar{K}$, then it is strictly optimal for a borrower with a good signal, i.e. $m = 0$, to repay. By continuity,

[28] To recast in the language of information design, it may be worth clarifying that the lender is recommended to take a *random* action when the borrower's history is B, rather than the information system recommending each pure action to the lender with positive probability.

repayment is also optimal for a borrower with $m = 0$ for an interval of values, $p \in [0, \tilde{p}]$, where $\tilde{p} > 0$ is the threshold where such a borrower is indifferent between repaying and defaulting. We restrict attention to values of p in this interval in what follows. Note that the best responses of a borrower with $m = 1$ are identical to those of a borrower with $m = 0$, for any p, since their continuation values are identical. Also, any increase in p increases the attractiveness of defaulting, and so a borrower with $m > 1$ will continue to default when $p > 0$.

The value function for a borrower with a good signal, i.e. $m = 0$, is given by:

$$\tilde{V}^{\bar{K}}(0,p) = (1-\delta) + \delta \left[\lambda \tilde{V}^{\bar{K}}(\bar{K},p) + (1-\lambda)\tilde{V}^{\bar{K}}(0,p) \right], \tag{10}$$

while the value function of a borrower with signal B is given by

$$\tilde{V}^{\bar{K}}(m,p) = \begin{cases} p(1-\delta) + \delta \left[p\lambda V^{\bar{K}}(\bar{K},p) + (1-p\lambda)V^{\bar{K}}(m-1,p) \right] & \text{if } m = 1, \\ p(1-\delta)(1+g) + \delta \left[pV^{\bar{K}}(\bar{K},p) + (1-p)V^{\bar{K}}(m-1,p) \right] & \text{if } m > 1. \end{cases} \tag{11}$$

Any $p \in [0, \tilde{p}]$, in conjunction with the borrower responses and exogenous default probability λ, induces a unique invariant distribution μ on the state space $\{0, 1, 2, \ldots, \bar{K}\}$. A borrower with $m > 1$ transits to $m - 1$ if she does not get a loan, and to $m = \bar{K}$ if she does get a loan, and thus

$$\mu_{m-1} = (1-p)\mu_m \quad \text{if } m > 1. \tag{12}$$

The measure $\mu_{\bar{K}}$ equals both the inflow of involuntary defaulters, who defect at rate λ, and the inflow of deliberate defaulters from states $m > 1$, so that

$$\mu_{\bar{K}} = \lambda \left(\mu_0 + p\mu_1 \right) + p \sum_{m=2}^{\bar{K}} \mu_m. \tag{13}$$

Finally, a borrower with $m = 1$ transits to $m = 0$ unless she gets a loan and suffers involuntary default. Thus, the measure of agents with $m = 0$, i.e. with a good credit history, equals

$$\mu_0 = (1-\lambda)\mu_0 + (1-p\lambda)\mu_1. \tag{14}$$

Since μ_m depends on p and also on the repayment probability for borrowers with types $m \in \{0, 1\}$, which equals 1, we write it henceforth as $\mu_m(p, 1)$. Figure 4 depicts the invariant distribution over the values of $m \in \{1, 2.., \bar{K}\}$, conditional on signal B, for two values of p. The horizontal line depicts the conditional distribution when $p = 0$, which is uniform. The

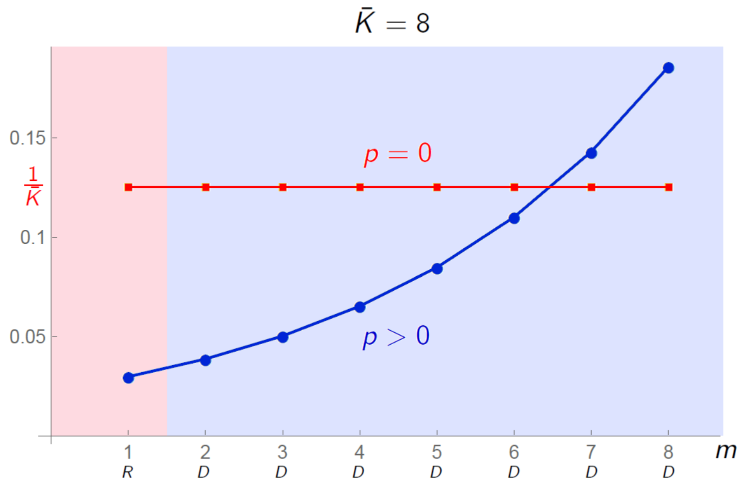

Figure 4: *Stationary probabilities conditional on signal* B*:* $\mu_m/(1-\mu_0)$ *for* $m = 1, \ldots, \bar{K}$*. Illustrated for* $p = 0$ *and* $p = 0.23$*.*

distribution conditional on $p > 0$ is upward sloping since higher values of p increase $\mu_{\bar{K}}$ and depress μ_1.

The probability that a loan made at history B is repaid is

$$\pi(p, 1) := \frac{\mu_1(p, 1)}{1 - \mu_0(p, 1)}. \tag{15}$$

In Appendix A.6.1 we show that this is a continuous and strictly decreasing function of p. Intuitively, higher values of p result in more defaults at B, increasing the slope of the conditional distribution. Thus if $\pi(\tilde{p}, 1) \leq \frac{\ell}{1+\ell}$, the intermediate value theorem implies that there exists a value of $p \in (0, \tilde{p}]$ such that $\pi(p, 1) = \frac{\ell}{1+\ell}$. This proves the existence of a mixed strategy equilibrium where all borrowers have pure best responses.

If loans are so profitable that $\pi(\tilde{p}, 1) > \frac{\ell}{1+\ell}$, then an equilibrium also requires mixing by the borrower. At $\tilde{p}$, the borrower with $m = 1$ is indifferent between repaying and defaulting on a loan. In this case, a borrower with a good signal (i.e. with $m = 0$) is also indifferent between repaying and defaulting, and there is a continuum of equilibria where these two types repay with different probabilities. However, only the equilibrium in which both types, $m = 1$ and $m = 0$, repay with the same probability, q, is purifiable. We focus our analysis on this equilibrium.

Let $\mu(\tilde{p}, q)$ denote the invariant distribution over values of m induced by this strategy

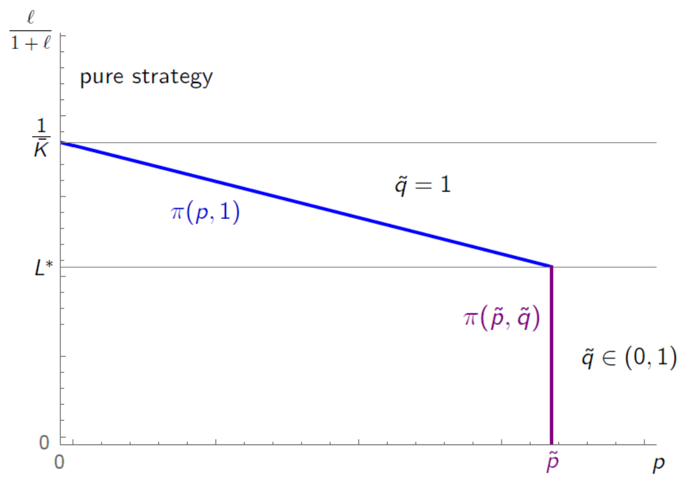

Figure 5: *The functions $\pi(p, 1)$ and $\pi(\tilde{p}, q)$. Let $L^* = \frac{\ell^*}{1+\ell^*}$. When $L^* \leq \ell/(1+\ell) \leq 1/\bar{K}$ there exists a mixed strategy equilibrium in which all types of borrowers use a pure strategy. When $0 < \ell/(1+\ell) < L^*$, there exists a mixed strategy equilibrium in which types $m = 0$ and $m = 1$ repay with probability q setting $\pi(\tilde{p}, q) = \ell/(1+\ell)$.*

profile. The probability that a loan made at history B is repaid is now

$$\pi(\tilde{p}, q) := \frac{q\,\mu_1(\tilde{p}, q)}{1 - \mu_0(\tilde{p}, q)} = q\,\pi(\tilde{p}, 1). \tag{16}$$

We establish the second equality in Appendix A.6.2. Clearly, $\pi(\tilde{p}, q)$ is a continuous and strictly increasing function of q. Since we are considering the case where $\pi(\tilde{p}, 1) > \frac{\ell}{1+\ell}$, and since $\pi(\tilde{p}, 0) = 0$, the intermediate value theorem implies that there exists a value of q setting the repayment probability $\pi(\tilde{p}, q)$ equal to $\frac{\ell}{1+\ell}$, so that lenders are indifferent between lending and not lending to a borrower with signal B. Figure 5 illustrates our analysis, with p on the horizontal axis and the repayment probability after B, $\pi(p, q)$, on the vertical axis. When $q = 1$, so that a borrower with $m = 1$ repays for sure, $\pi(p, 1)$ is given by the downward sloping (blue) line. This stops at $\tilde{p}$, and further declines of the repayment probability are achieved by reducing q, along the (purple) vertical line. Figure 6 illustrates the corresponding equilibrium payoff for the borrower and the lender. We have therefore established the following proposition.

Proposition 4 *Suppose $\bar{K} \geq 2$. If $0 < \ell < \frac{1}{\bar{K}-1}$, there exists a purifiable mixed equilibrium under the simple information structure and $\bar{K}$ memory, where the borrower's payoff is strictly greater than $\bar{V}$. If ℓ is strictly greater than a threshold value ℓ^*, then the borrower plays a pure*

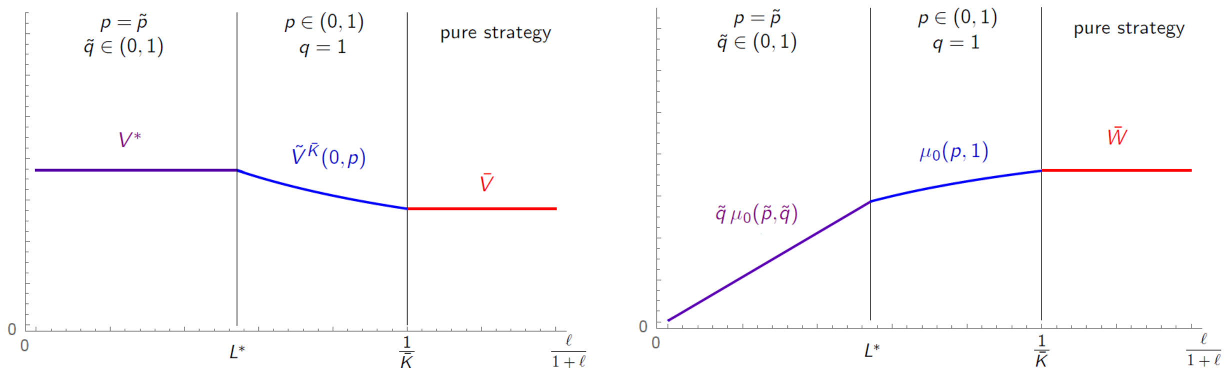

(a) *Borrower's equilibrium payoff.* (b) *Lender's equilibrium payoff.*

Figure 6: *Equilibrium payoffs under the simple information structure with $\bar{K}$ period memory.*

strategy, where she repays if $m \in \{0,1\}$. If $\ell \in (0,\ell^*]$, then loans are made with probability $\tilde{p}$ after B, and borrower types $m \in \{0,1\}$ repay with probability $\tilde{q}$ so as to make the lender indifferent between lending and not lending at B.

The borrower payoff in the mixed equilibrium lies between $\bar{V}$ and V^*. It is strictly greater than $\bar{V}$ since the exclusion for $\bar{K}$ periods is random, and she gets a loan with positive probability. When ℓ is so low that the borrower also mixes, then she gets the payoff V^* — this follows from the fact that her incentive constraint is satisfied with equality when she has a good signal. Thus, when lending becomes more profitable, the borrower's payoff increases in the mixed equilibrium. However, payoffs for the lender are strictly less than $\bar{W}$. Since the lender only makes positive profits when she lends to a borrower with a good signal, steady state profits equal the proportion of borrowers with signal G. This proportion falls as p increases. When the borrower also mixes, and defaults after a good signal with some probability, the lender's profits fall further, and as ℓ tends to zero, so do profits.

Appendix A.6.3 proves that the equilibria described in the proposition are purifiable. In the case where both lenders and borrowers with $m \in \{0,1\}$ mix, it is worth noting that the equilibrium is not *regular* — it is not isolated, since it is one of a continuum of equilibria.[29] Here, we give a brief intuition for why purification selects a unique mixed equilibrium. When the lending probability after B equals $\tilde{p}$, and borrower types $m \in \{0,1\}$ are indifferent between defaulting and repaying, there is a continuum of equilibria in the unperturbed game. The borrower repayment probability $\tilde{q}$ at $m=1$ is pinned down by the

[29]The standard proofs for purifiability, in Harsanyi (1973), Govindan, Reny, and Robson (2003) and Doraszelski and Escobar (2010), apply for regular equilibria.

equilibrium condition $\pi(\tilde{p}, \tilde{q}) = \frac{\ell}{1+\ell}$, but the repayment probability at $m = 0$ can take any value in some interval. In particular, one can have the borrower repaying with probability one at $m = 1$, but with probability $\tilde{q}$ at $m = 0$. However, the difference between borrower histories at these two values of m is payoff-irrelevant for the borrower, since no future lender can see this difference. The payoff shocks prevent any such conditioning, and only one of these equilibria can be purified, namely the one where the repayment probability at $m = 0$ also equals $\tilde{q}$. Thus, although the equilibrium where the borrower repays for sure when $m = 1$ is better for the lenders (and no worse for the borrowers), it cannot be sustained.

An Example: We consider parameter values such that $\bar{K} = 4$.[30]. Consider the pure strategy profile when $K = \bar{K}$, where the lender extends a loan only after G. From our previous analysis, the invariant distribution over $\{1, \ldots, K\}$ is uniform, and the probability that the lender is repaid if lending at B equals $\frac{1}{4}$. So if $\ell > \frac{1}{3}$, a pure strategy equilibrium exists. The expected payoff to a borrower with a good history is $\bar{V} = 0.763$. The lender's expected payoff equals the probability of encountering a borrower with a good history, which is $\bar{W} = 0.714$.

If $\ell < \frac{1}{3}$, lending after B is too profitable and a pure strategy equilibrium with $\bar{K}$-period memory does not exist. A pure strategy equilibrium with longer memory exists, but can be very inefficient. For example, if $\ell = 0.315$, we need $K = 30$, in which case $m^*(K) = 7$ and the lender strictly prefers not lending at B. The invariant distribution over borrower types has only a quarter of the population with a clean history, so that the lender's payoff equals 0.25, strictly less than $\bar{W}$. Since exclusion is very long, the borrower's payoff at a clean history is also low, at 0.537, which is strictly less than $\bar{V}$.

In the mixed strategy equilibrium with 4-period memory, the lender offers a loan with probability $\tilde{p} = 0.028$ to a borrower on observing B. This equilibrium is considerably more efficient than the pure equilibrium with 30-period memory. The proportion of borrowers with clean histories is $\mu_0 = 0.705$, and this is also the lender's expected payoff. The expected payoff to a borrower with a clean history is $0.775 > \bar{V}$, since she sometimes gets a loan even when the signal is B.

If ℓ is smaller, say 0.1, then the mixed equilibrium also requires random repayment by borrowers of types $m \in \{0, 1\}$.[31] The lending probability after B is $\tilde{p} = 0.034$, and the repayment probability is $\tilde{q} = 0.383$ for $m = 0$ and $m = 1$. The lender's payoff in this

[30] Specifically, $\delta = 0.9, \lambda = 0.1$ and $g = 2$.

[31] If $\ell = 0.1$, $K = 89$ periods of exclusion are needed to support a pure strategy equilibrium. In this case $m^*(K) = 8$. The lender's payoff is $\mu_0 = 0.1$ and the borrower's payoff at a clean history is 0.526.

equilibrium is substantially lower: $\mu_0(\tilde{q} - \ell(1-\tilde{q})) = 0.084$. This is largely because a lower fraction of the population has a good history: $\mu_0 = 0.262$. The payoff to the borrower with a clean history equals $V^* = 0.778$.

At the same value of ℓ, there exists another equilibrium at which the lending probability at G is one and at B is $\tilde{p} = 0.034$, and where the borrower with $m = 0$ repays the loan with certainty, while the borrower with $m = 1$ repays it with probability $q_1 = 0.383$. The lender's payoff in this equilibrium, $\mu_0 = 0.699$, is substantially higher than in the equilibrium where both $m = 0$ and $m = 1$ mix with the same probability. The borrower's payoff at $m = 0$ remains V^*. Notice that this equilibrium is not purifiable, but it Pareto-dominates the purifiable equilibrium at which $m = 0$ and $m = 1$ repay a loan with the same probability. □

Taking stock: we have demonstrated that for any parameter values, the simple information structure can ensure a borrower payoff of at least $\bar{V}$, no matter how profitable loans are (or, equivalently, how small ℓ is). When ℓ is large, a pure strategy equilibrium exists, and yields borrower payoffs of $\bar{V}$ and lender payoffs of $\bar{W}$. We can also achieve fully efficient payoffs V^* and W^* with a random signal structure. When ℓ is small, the mixed strategy equilibrium yields borrower payoffs greater than $\bar{V}$, but lender payoffs strictly below $\bar{W}$. For very small values of ℓ, lender payoffs can be extremely low.

The next sections examine how outcomes may be improved for lenders without compromising on achieving a borrower payoff of $\bar{V}$. In Section 6 we show that a non-monotone information structure can be used to support a pure strategy equilibrium with minimal borrower exclusion. In Section 7.2 we show that preventing lenders from chasing borrowers, by requiring initial applications from the latter, can also do the same under the simple information structure. In both cases, pure strategy equilibria achieve the constrained efficient payoffs $\bar{V}$ and $\bar{W}$.

6 Non-monotone Information

How can we discipline lenders when ℓ is small? One idea is as follows: suppose that we reward borrowers for defaulting on a lender who should not have made a loan. This makes it more likely that a borrower with a bad credit history will default, and dissuades lending to them. We can implement this idea, under bounded memory, by assigning a good credit history to borrowers with *two* defaults, while those with a single default are assigned a bad credit history. We call such an information structure *non-monotone*, since borrowers with no defaults are pooled with borrowers with two defaults, while one-default borrowers are excluded from this pool.

There still remains the problem identified in Section 3.2, that the incentives of a borrower with a single default which occurred exactly K periods ago (i.e. a borrower with $m = 1$) are identical to that of a borrower with no defaults. Thus in any purifiable equilibrium, the behavior of the two borrowers must be identical, and if the borrower with no defaults repays for sure, so must the borrower with a single default and $m = 1$. This suggests that the probability of repayment of a lender with a bad credit history cannot be reduced below $\frac{1}{K}$. Nonetheless, the following proposition constructs an information structure where this repayment probability is zero. The trick here is to allow for memory that is one period longer than is required for borrower incentives, so that $K = \bar{K} + 1$. The additional period is not used to punish the borrower, but instead to retain information for disciplining the lender.

Proposition 5 *For any $\ell > 0$, there exists a non-monotone informational structure, and a sequentially strict equilibrium that achieves borrower payoff $\bar{V}$ and lender payoff $\bar{W}$.*

Proof. Let the length of memory be $K = \bar{K} + 1$, and let $N_{\mathcal{D}}$ denote the number of instances of $\mathcal{D}$ in the last K periods. The lender observes credit history G if $N_{\mathcal{D}} \in \{0, 2\}$, or if $N_{\mathcal{D}} = 1$ and $m = 1$. Otherwise, the credit history is B — in particular, if $N_{\mathcal{D}} = 1$ and $m > 1$. The lender lends after G and does not lend after B. Consequently, a defaulting lender with no instance of $\mathcal{D}$ in the past is excluded for $\bar{K}$ periods, and thus repayment is optimal on being given a loan. Consider a borrower with $N_{\mathcal{D}} = 1$ and $m > 1$, who has credit history B. If she receives a loan, this borrower will default, since by doing so she gets a good credit history in the next period (since $N_{\mathcal{D}} = 2$ in the next period). Thus defaulting raises her continuation value as well as current payoff, relative to repayment. Consequently, the probability of repayment of a loan made to a borrower with history B is zero, and thus for any $\ell > 0$, lending at B is strictly unprofitable. Since exclusion is for $\bar{K}$ periods, the payoffs are as stated in the proposition. ■

Remark 3 *The proposition applies also when $\bar{K} = 1$ — in contrast with Proposition 3, which does not.*

The novelty of the information structure here is that it dissuades lenders from lending to defaulters while respecting the bounded memory constraint. As we have seen in Section 3.1, this is easy to achieve with infinite memory. With bounded memory, although a non-monotone information structure allows us to achieve the payoffs $\bar{V}$ and $\bar{W}$, the payoffs V^*

and W^* that are not subject to the integer constraint remain out of reach. Indeed, a random version of the non-monotone information structure turns out to violate incentives.

Non-monotonicity of the information is an unrealistic feature, since a borrower with a worse default record is given a better rating than one with a single default. Furthermore, it may be vulnerable to manipulation, if we take into account real-world considerations that are not explicitly modelled. A borrower with two defaults in the last $K-1$ periods is ensured of a continuation value corresponding to being able to default without consequence twice in every K periods. This is very attractive, and a borrower with a single default may be willing to pay a large bribe to a lender, in exchange for the privilege of defaulting a second time.

7 More General Games or Outcomes

We have fixed a stage game, Γ, where each player moves at most once, and considered the sustainability of an arbitrary Pareto-efficient outcome in this game. We now show that, when the stage game has multiple Pareto-efficient outcomes, higher payoffs for both players might be achieved by supporting an outcome that is more favorable for the player who is subject to moral hazard. Payoffs may also be improved if the player who is subject to moral hazard has to initiate the interaction — with our coarse information structure, such a player has private information about her previous history, so that her initial move signals her private history. Finally, we show how our analysis may be extended to more general stage games, as long as moral hazard is effectively one-sided. Throughout this section, we restrict attention to binary information structures that are monotone.

7.1 Multiple Efficient Profiles

How is the analysis affected when making the terms of trade more favourable for player 2 — the one player subject to moral hazard — at the expense of player 1? In the lender-borrower interaction, this would correspond to the lender extending a loan with a smaller interest rate. In the seller-buyer example, this corresponds to the buyer purchasing a higher-priced product that is no more costly for the seller to produce.

Varying terms of trade corresponds to multiple Pareto-efficient outcomes in the stage game, as noted in Section 2.1. To illustrate, consider the augmented trust game whose extensive form is depicted in Figure 7a. Player 1 has two trusting actions, Y and $\tilde{Y}$, in addition to his outside option, N. After each of the trusting actions, player 2 must choose whether to repay or default. With probability λ, player 2 has no choice, and must default

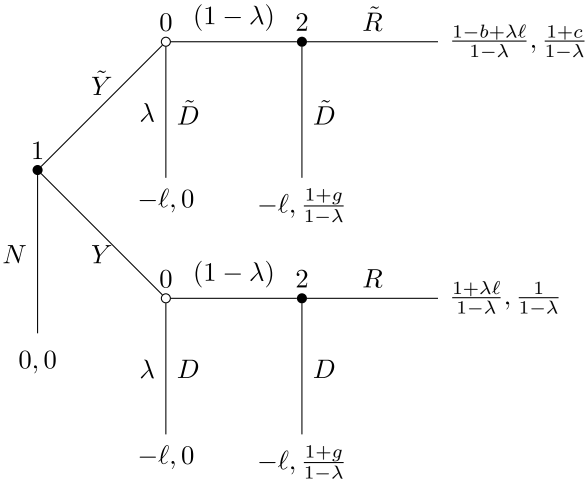

(a) *Extensive Form.*

	$\tilde{D}, R$	$\tilde{D}, D$
$\tilde{Y}$	$1-b, 1+c$	$-\ell, 1+g$
N	$0, 0$	$0, 0$

	$\tilde{R}, D$	$\tilde{D}, D$
Y	$1, 1$	$-\ell, 1+g$
N	$0, 0$	$0, 0$

(b) *Reduced strategic forms.*

Figure 7: *Extensive and reduced strategic forms of the augmented trust game*

involuntarily. We assume $b \in (0, 1)$ and $c \in (0, g)$, so that the outcomes corresponding to the action profiles $(\tilde{Y}, \tilde{R})$ and (Y, R) are both Pareto-efficient.

Define the stage game strategy profile σ^* by Y for player 1 and $(\tilde{D}, R)$ for player 2 — this yields the payoff vector $(1, 1)$. The reduced strategic form that is relevant for supporting the corresponding outcome is depicted in the bottom panel of Figure 7b, where the two strategies for each player are σ_i^* and the backward induction strategy, $\bar{\sigma}_i$. Define the stage game strategy profile $\tilde{\sigma}^*$ by $\tilde{Y}$ for player 1 and $(\tilde{R}, D)$ for player 2 — this yields the payoff vector $(1 - b, 1 + c)$. The reduced strategic form that is relevant for supporting this outcome is depicted in the top panel of Figure 7b, where the two strategies for each player are $\tilde{\sigma}_i^*$ and the backward induction strategy, $\bar{\sigma}_i$. Observe that distinct reduced strategic forms are relevant depending on the outcome that we seek to support.

Consider the sustainability of the outcome of the profile $(\tilde{Y}, \tilde{R})$. Player 2 is subject to moral hazard, but her moral hazard problem is less acute than under the profile (Y, R) — the terms of repayment are less onerous under $(\tilde{Y}, \tilde{R})$, since $0 < c < g$. Consequently, the length of punishment, $\tilde{K}$, required to support $(\tilde{Y}, \tilde{R})$ is reduced as compared to the length of punishment, $\bar{K}$, required for sustaining (Y, R). In addition, the profitability of a loan for the lender is also reduced, and so the condition for the existence of a pure strategy equilibrium with punishments of minimal length becomes:[32]

$$\frac{\ell}{1-x} > \frac{1}{\tilde{K}-1}. \tag{17}$$

Recall that the outcome (Y, R) can be supported via a pure strategy equilibrium when $\ell > \frac{1}{\bar{K}-1}$. Thus, supporting the borrower-optimal outcome $(\tilde{Y}, \tilde{R})$ via pure strategies is feasible for a larger set of parameter values than supporting the lender-optimal outcome (Y, R) via pure strategies.

Supporting the borrower-optimal outcome is clearly better for the borrower — she gets better loan terms, and the period of exclusion following defaults is reduced. In some cases, it may also be better for the lender, particularly if the lender-optimal outcome can only be supported under the simple information structure via mixed strategies, since the mixed equilibrium yields low payoffs for the lender.

To summarize: supporting an outcome that is better for the player who is subject to moral hazard (player 2), always improves the equilibrium payoff of this player. It may also improve the payoff of the player who is not subject to moral hazard (player 1) by not requiring the use of mixed strategies. The question remains, how can we ensure that player 1 chooses action $\tilde{Y}$ rather than action Y? As explained in Section 2.2, it suffices to pool stage-game outcomes that arise after player 1 chooses Y with the outcome N, so that, for third parties, a default on a loan with unfavourable terms is indistinguishable from not having received a loan. It may be argued that the information system may not plausibly distinguish defaults according to the terms of the loam, and that all defaults must be treated in the same way. Fortunately, a coarse information structure that distinguishes only three subsets of stage game outcomes, "no loan", "default" and "repayment", where the last two sets include all default (resp. repayment) irrespective of loan terms[33], can still deter the lender from offering worse terms to the borrower. If $\tilde{K} < \bar{K}$, then a punishment length of $\tilde{K}$ suffices to deter defaults after $\tilde{Y}$ but not deter defaults after Y. This disciplines the lender and dissuades him from offering more onerous terms.[34] This illustrates the merit of an information system that does not punish the borrower for too long — it deters the lender from demanding more

[32] This condition can be derived from our previous results by re-normalizing payoffs, so that the payoff from $(\tilde{Y}, \tilde{R})$ equals $(1, 1)$.

[33] That is, an information structure based on the partition $\{\mathcal{P}_1, \mathcal{P}_2, \mathcal{P}_3\}$, where $\mathcal{P}_1 := \{(N)\}$, $\mathcal{P}_2 := \{(Y, R), (\tilde{Y}, \tilde{R})\}$, and $\mathcal{P}_3 := \{(Y, D, 0), (Y, D, 2), (\tilde{Y}, \tilde{D}, 0), (\tilde{Y}, \tilde{D}, 2)\}$

[34] Even if $\tilde{K} = \bar{K}$, the required length of punishment without integer constraints is strictly lower for supporting $(\tilde{Y}, \tilde{R})$, so that there exists a random information structure that deters defaults after $\tilde{Y}$ but not after Y.

favorable terms, and may thereby cause a Pareto-improvement that also benefits the lender.

7.2 Costly initiation by the player who must be trusted

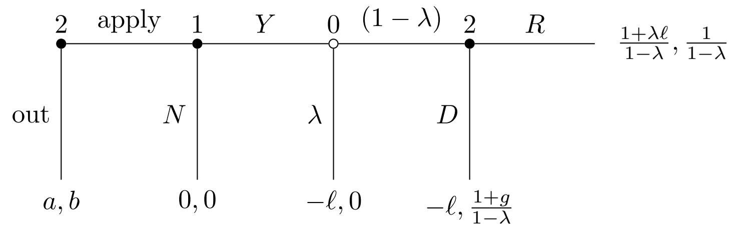

Figure 8: *Trust game with costly initiation.*

Our analysis so far has assumed a game where each player moves at most once. Suppose now that the game is more complex, with the first move made by the player who needs to be trusted subsequently. That is, the trust game is augmented by a prior stage, where player 2, the player subject to moral hazard, has to initiate the interaction. We model this via the extensive form game set in Figure 8. First, player 2 (the borrower) must decide whether to initiate the interaction with player 1, the lender, at a small cost $a > 0$, i.e. she must decide whether or not to "apply". If she does not apply, the game ends, with a payoff $b > 0$ for player 2. If she applies, the trust game is played. The backwards induction outcome is *out*, and the borrower does not apply. We assume the simple information structure with $\bar{K}$ periods memory.[35]

Under this information structure, requiring prior application transforms the interaction between an individual lender and borrower into a signaling game. Recall that a borrower with credit history B has private information regarding m, the number of periods without default that must elapse before her credit history becomes good. Since the informed party moves first, her application decision signals her private information. Consider the following strategy profile, σ^*. Borrowers with history G apply, are given a loan and repay. Borrowers of type B do not apply; if they do make an application, the lender rejects their application; if the lender accepts their application, the borrower defaults if $m > m^*$ and repays if $m \leq m^*$. The following proposition shows that if $K = \bar{K} \geq 2$, then this strategy profile is a sequentially

[35]For the full efficiency result, we use the random version of the simple information structure, under which the borrower gets a good signal either if she has not defaulted in the last $\bar{K}$ periods, and has defaulted exactly once $\bar{K}$ periods ago, in which case she gets signal G with probability x.

strict equilibrium, where the lender believes that an applicant with signal B will default with probability one, and these beliefs are implied by the D1 criterion of Cho and Kreps (1987). In other words, punishments are of minimal length, no matter how profitable loans are.

Proposition 6 *Assume that $K = \bar{K} \geq 2$ and the random version of the simple information structure. For any $\ell > 0$, σ^* is a sequentially strict perfect Bayesian equilibrium, that approximates payoffs V^* and W^*, with borrower beliefs satisfying the D1 criterion that assigns probability one to an applicant with a bad credit history defaulting.*

Proof. The borrower's strategy is a strict best response to the lender's strategy, since application is costly ($a > 0$). Given the lender's beliefs, his strategy is (strictly) sequentially rational. It remains to verify that the beliefs of the lender are implied by the D1 criterion. Suppose that a borrower with a bad credit history applies for a loan, and gets it. Since $m^* = 1$, only this type of borrower will repay, (cf. Section 4.1). Consider a mixed response of the lender to a loan application, whereby he gives a loan with probability q on observing a B applicant, where q is chosen so that the type of applicant who intends to repay (i.e. one with $m = m^* = 1$) is indifferent between applying or not. Thus, q satisfies

$$(1-\delta)a = q\big[1 - \delta + \delta\lambda\big(V(K) - V(0)\big)\big]. \tag{18}$$

Now consider a borrower of type $m' > 1$, whose optimal strategy is to default on the loan, if she receives it. Since default is optimal, her net benefit from applying, relative to not applying, equals

$$q\big[(1-\delta)(1+g) + \delta\big(V(K) - V(m'-1)\big)\big] - (1-\delta)a. \tag{19}$$

We now show that expression (19) above is strictly positive. Substituting for $(1-\delta)a$ from (18), and dividing by q, we see that the sign of (19) is the same as that of

$$(1-\delta)g + \delta(1-\lambda)\big(V(K) - V(m'-1)\big) + \delta\lambda\big(V(0) - V(m'-1)\big). \tag{20}$$

Since default is optimal for type m',

$$(1-\delta)g + \delta(1-\lambda)[V(K) - V(m'-1)] > 0,$$

establishing that the sum of the first two terms in (20) is strictly positive. Also, $V(0) > V(m')$, and so the third term as well as the overall expression in (20) is strictly positive.

Thus, type m' has a strict incentive to apply whenever m^* is indifferent.

We conclude therefore that a borrower who intends to default strictly prefers to apply, if q is such that any type of borrower who does not intend to default is indifferent. Thus the D1 criterion implies that in the equilibrium σ^*, the lender must assign probability one to defaulting types when he sees an application from a borrower with a B credit history. ■

More generally, consider a two-player game Γ where player 2 has an incentive to deviate from the path to the Pareto-efficient outcome, but player 1 does not. Augment this game by allowing player 2 an initial choice of actions, one of which leads to the game Γ, while each of the others leads to a sequential move continuation game. For generic payoffs in the new game, either the unique backwards induction outcome has player 2 choosing the action that leads to Γ, in which case the analysis is unaffected. Or, it has player 2 choosing a different action, which yields a greater payoff for player 2 than the backwards induction payoff in Γ. In the latter case, the above proposition applies: one gets full efficiency with a simple, binary information structure, and the beliefs of player 1 that support this equilibrium satisfy the D1 criterion.

We conclude that the party who must be trusted has to initiate the interaction, and when it somewhat costly to initiate an interaction which is not reciprocated, they efficient payoffs can be sustained for both parties. In the resulting signaling game, the fully efficient payoffs V^* and W^* are equilibrium payoffs.

7.3 General Sequential-Move Games

Consider a generic two-player game of perfect information and no moves by nature, and a Pareto-efficient outcome, y^*, that strictly Pareto-dominates the backwards induction outcome $\bar{y}$. Observe that, in the absence of moves by nature, the outcome y^* can be identified with a unique terminal node, z^*, and similarly, $\bar{y}$ with a unique $\bar{z}$. Let both players conjecture that any deviation from the path to z^* is followed by players continuing with the backward induction strategies. Suppose now that only one player has an incentive to deviate from the path to z^* given this conjecture about continuation play, and this incentive arises at a single node. The analysis of this paper applies to the sustainability of z^* any such game.

To illustrate how general these results are, consider the centipede game in Figure 9, where the backwards induction strategy profile has players choosing *down* (d_t) at every history. Consider the Pareto-efficient terminal node z^* that is reached when players choose *right* (r_t) at every history. *If players expect the outcome* z^*, then only one player has an incentive to

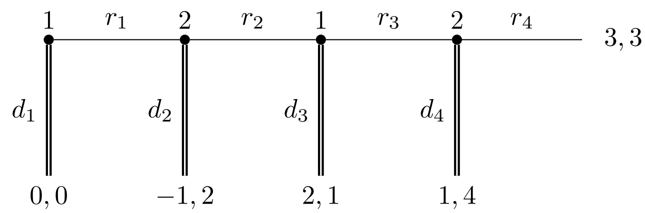

Figure 9: *Centipede game.*

deviate along the path to z^*: player 2. Furthermore, there is only a single decision node at which 2 has an incentive to deviate: the last one.

Now suppose that this game is played repeatedly in a random matching environment. Consider an information structure where, before the game is played in a match, player 1 is provided coarse information on the last K outcomes of player 2. Specifically, signal B is observed if and only if 2 has played d_4 in at least one of the last K periods; signal G is observed otherwise. Suppose that 1 plays d_1 upon observing signal B, and let $\bar{K}$ denote the smallest punishment length that suffices to incentivize player 2 to play r_4 at G.[36] As long as $K \geq \max\{\bar{K}, 2\}$, there exists an equilibrium that supports the efficient outcome z^* under this information structure. Furthermore, it is costly for player 2 to choose r_2 at her first decision node if she expects player 1 to continue with his backwards induction action d_3. This ensures that if 2 plays r_2 and has a bad signal, then the D1 criterion implies that, at his second information set, 1 believes that if he plays r_3, 2 will continue with d_4.

In Appendix A.7, we show that this argument generalizes. Fix a Pareto-efficient terminal node z^* that strictly Pareto-dominates the backwards induction outcome in a two-player game of perfect information. Suppose that:

- Only one player, $\hat{i}$, has an incentive to deviate from the path to z^*.
- This incentive to deviate exists only at a single node, $\hat{x}$.
- Player $\hat{i}$ has a decision node x that precedes $\hat{x}$ on the path to z^* where her backwards induction strategy prescribes deviating from $\phi(z^*)$ at x.

Under these assumptions, there exists a simple, binary information structure and an equilibrium that sustains play of the efficient outcome z^*.[37] In this equilibrium, players play the

[36] If g denotes player 2's payoff gain from playing d_4 instead of r_4, then $\bar{K}$ is the smallest integer K that satisfies the inequality $g \leq \frac{\delta(1-\delta^K)}{1-\delta}$, where δ is player 2's discount factor.

[37] In fact, our analysis applies to stage games of perfect information with moves of nature and an arbitrary

outcome z^* when the signal about player $\hat{i}$ is good, and they play the backwards induction outcome when the signal is bad. Players' beliefs satisfy the D1 criterion, and it suffices to pick punishments of a minimum length that is sufficient to discipline player $\hat{i}$.

If the last of the above assumptions is not satisfied, then one also needs to ensure that the player who does not have an incentive to deviate from the path to z^* does indeed punish player $\hat{i}$ when he sees a bad signal regarding her. This concern has been the focus of most of this paper, and methods similar to those set out in the context of the lender-borrower game can be used. To avoid repetition, we do not present the details here.

If the first of these assumptions is not satisfied, and both players have an incentive to deviate from the path to z^*, then providing incentives for both players becomes more difficult. This would be the case, for example, if the payoff of player 1 after r_3 was changed to 4. As we will see in Section 7.4, it is difficult to provide incentives for a player to condition his behavior on a signal regarding his opponent via *future play.* Incentives for such conditioning have to be provided within the period. (In the lender-borrower game, the lender has only within-period incentives to condition his behavior on the borrower's signal). This makes the analysis of this case more difficult, and we leave this for future work.

7.4 The Irrelevance of Information on Lenders

We have assumed that there is information on the past outcomes of player 2, the one who is subject to moral hazard, but there is no information on player 1. Nonetheless, in our lender-borrower example, disciplining lenders to not lend to borrowers in bad standing is a major problem. Can information on lender behavior be used to prevent lenders from making loans to borrowers who have recently defaulted? Suppose that the borrower also observes information on the past outcomes of the lender. It is easy to see that this additional information is not useful, since observing the outcome in any interaction of the lender does not convey information on whether the lender should have made the loan or not in the first place.

Now, suppose that the borrower not only observes the outcomes in the lender's past interactions, but also the information that the lender received about the borrowers he interacted with. For example, if the lender's information regarding a borrower is a binary signal, then the borrower today observes the outcome in each past K interaction plus the K realizations of the binary signal observed by the lender. In particular, if the lender lent yesterday to

finite number of players, provided that these conditions hold, but we prove this only for the case of two players in games without moves of nature.

a borrower with a bad signal, the borrower today can see this. The question is: can we leverage this additional information in order to discipline the lender, so that he does not lend to borrowers with a bad signal?

Unfortunately, the answer is no. Consider a borrower's repayment decision conditional on obtaining a loan. While the borrower can observe information on the lender's past behavior, in any purifiable equilibrium, she will not condition her behavior on this information. Since no future lender will observe this information, but will only observe the outcome in this and past interactions of the borrower, information on the lender's past behavior is payoff-irrelevant.

Thus in the perturbed game where the borrower is subject to payoff shocks, the borrower can play differently after two different lender histories, h_1 and h'_1, only for a set of payoff shocks that have Lebesgue measure zero. Hence, in any purifiable equilibrium of the unperturbed game, a borrower cannot condition her repayment behavior on any information regarding the lender. So our assumption that the borrower observes no information regarding the lender is without loss of generality.

This discussion also illuminates the interplay between information and incentives that underlies our analysis. Since a lender's future continuation value cannot depend on his current behavior, incentives have to be provided *within the period*. This is possible, since borrowers with a bad signal have a higher incentive to default than those with a good signal. This also explains why our results extend to general games where only one player has an incentive to deviate from the path to an efficient outcome, but not to games where both players have an incentive to deviate.

Finally, one might ask whether our analysis is robust to different forms of bounded memory. For example, suppose that loans that are defaulted on are expunged from the record, but otherwise, there is no uniform bound on memory. In Appendix A.8, we show that would get similar results under this formulation.

8 Conclusion

Existing work on community enforcement has focused entirely on simultaneous-move games, and mainly on the prisoner's dilemma. Many economic interactions have a sequential structure, the trust game being a paradigmatic example. Although moral hazard is one-sided, incentivizing the side that is not subject to moral hazard to exclude offenders turns out to be non-trivial. Our substantive results show that *endogenous adverse selection* can be leveraged by an information designer to provide such incentives.

Our analysis offers suggestions for further empirical work on consumer credit scores and consumer bankruptcy. A notable finding in existing empirical work is that credit scores increase discontinuously in the quarter when a bankruptcy "flag" is removed from a consumer's record, and that this leads to increased access to credit (see Section 1.1). Our paper raises the question: do lenders increase credit to consumers whose bankruptcy flags are on the verge of being removed? Existing empirical work does not address this question, its focus being on the before-after comparison. Our paper also suggests that an analysis of the welfare properties of the credit scoring system would be useful. The stated purpose of credit scoring is to provide lenders with predictions of the likelihood than any one given borrower might default (and other delinquency) — its purpose is not to sustain socially efficient outcomes. Our analysis suggests that the two objectives may well be in conflict. Of course, an empirical evaluation of the system needs to take into account adverse selection.

A Appendix: For Online Publication

A.1 Proof of Lemma 1

Player 2 has an incentive to deviate since σ_2^* differs from the unique backwards induction strategy. For player 1, the best possible deviation is to his backwards induction strategy, given that player 2 responds with her backwards induction strategy after any deviation. Since y^* strictly Pareto-dominates $\bar{y}$, this deviation is unprofitable.

A.2 Proof of Lemma 3

First, let us make more precise the notion of measurability with respect to $\mathcal{O}$ that is invoked in the lemma. Player 2's history at date t, $h^t = (z^s)_{s=1}^{t-1}$, is the sequence of terminal nodes reached in all her previous interactions. Consider two period-t histories of player 2, $h^t = (z^s)_{s=1}^{t-1}$ and $\tilde{h}^t = (\tilde{z}^s)_{s=1}^{t-1}$. Write $h^t \sim \tilde{h}^t$ if $\mathcal{O}(z^s) = \mathcal{O}(\tilde{z}^s)$ for all $s \in \{1, 2.., t-1\}$. The strategy s for player 1 in the repeated game is *measurable with respect to the partition* $\mathcal{O}$ if $s(h^t) = s(\tilde{h}^t)$ whenever $h^t \sim \tilde{h}^t$.

We now prove the lemma. Suppose player 1 has played $a_1 \neq a_1^*$. Since the outcome today will belong to $\mathcal{N}$ no matter what player 2 plays, it is optimal for 2 to maximize her payoff, and play her backwards-induction action, $\bar{\sigma}_2(a_1)$. Suppose player 1 has played a_1^*. If player 2 plays any action $\sigma^*(a_1^*) = a_2^*$, then the outcome today belongs to $\mathcal{R}$. If player 2 plays any action $a_2 \neq a_2^*$, then the outcome today belongs to $\mathcal{D}$, and the payoff-maximizing action within the set $A_2(a_1^*) \setminus a_2^*$ is $\bar{\sigma}_2(a_1^*)$. Since player 2 responds to any action $a_1 \neq a_1^*$ by playing her backwards-induction strategy, it is optimal for player 1 to play his backward induction action if he does not play a_1^*.

A.3 Proofs Related to Section 3.1

We derive the upper bound on the borrower's value in any equilibrium, when the borrower's mixed strategies are not observable, and there is imperfect monitoring, i.e. the lender's strategy is measurable with respect to the partition $\mathcal{O}$. Let V^* be the supremum value of the borrower's payoff in any equilibrium. Note that in any period when her value is near V^*, the borrower must receive a loan. For the lender to agree to lend, the borrower must repay with positive probability. Thus repaying for sure in the current period must be optimal.

Thus V^* satisfies the inequality:

$$V^* \leq (1-\delta) + \delta \left[\lambda\, V^P + (1-\lambda)\, V^*\right]. \tag{A.1}$$

The right-hand side above is an upper bound on the borrower's value when she chooses to repay, and is derived as follows. When the borrower is able to repay, she gets no more than V^* tomorrow, since this is the supremum value in any equilibrium. V^P denotes the borrower's value following involuntary default. Since repaying must optimal in the current period, we have the following necessary incentive constraint

$$(1-\delta)(1+g) + \delta\, V^P \leq (1-\delta) + \delta \left[\lambda\, V^P + (1-\lambda)\, V^*\right],$$

where the left-hand side is the borrower's payoff from voluntary default. Rearranging gives

$$(1-\delta)\, g \leq \delta\, (1-\lambda)\, [V^* - V^P]. \tag{A.2}$$

By substituting inequality A.2 in A.1, we get

$$V^* \leq 1 - \frac{\lambda}{1-\lambda}\, g. \tag{A.3}$$

To see that the above bound on V^* is indeed achievable, let the difference $V^* - V^P$ be such that A.2 holds with equality. It is straightforward to verify that this implies that A.1 holds with equality. This yields the expression for V^* in equation 4.

A.4 Proofs Related to Sections 4.1 and 5.1

First, we show that if $K \geq \bar{K}$, then a borrower with $m = K$ has a strict incentive to default, i.e. we establish the inequality below (a restatement of (6)):

$$g > (1-\lambda)\delta^K V^K(0).$$

Since both δ^K and $V^K(0)$ are strictly decreasing in K, it suffices to prove this for $K = \bar{K}$. Here, we prove a stronger result that implies the first part of Lemma A.1. For $K = \bar{K}$ and for any $m > 1$,

$$(1-\delta)g > \delta(1-\lambda)\left(V^{\bar{K}}(m-1) - V^{\bar{K}}(\bar{K})\right).$$

Since $V^{\bar{K}}(m)$ is strictly decreasing in m, it suffices to prove this for $m = 2$. That is, we need to establish the inequality

$$(1-\delta)g > \delta(1-\lambda)\left(V^{\bar{K}}(1) - V^{\bar{K}}(\bar{K})\right) = \delta(1-\lambda)(\delta - \delta^{\bar{K}})V^{\bar{K}}(0). \tag{A.4}$$

By the definition of $\bar{K}$, the incentive constraint (1) is not satisfied for $K = \bar{K} - 1$, so that

$$(1-\delta)g > \delta(1-\lambda)\left(V^{\bar{K}-1}(0) - V^{\bar{K}-1}(\bar{K}-1)\right) = \delta(1-\lambda)(1 - \delta^{\bar{K}-1})V^{\bar{K}-1}(0).$$

Thus, to prove (A.4), it suffices to show that

$$V^{\bar{K}-1}(0) > \delta\, V^{\bar{K}}(0).$$

Since $V^{\bar{K}-1}(0) > V^{\bar{K}}(0)$, the above inequality is proved. We have therefore established (6).

We now solve for $m^{\dagger}$, the real value of m that sets (7) equal to zero:

$$m^{\dagger}(K) = \ln\left[\frac{(1-\delta)g}{(1-\lambda)V^{K}(0)} + \delta^{K+1}\right] / \ln\delta.$$

Lemma A.1 *If $K = \bar{K}$, then $m^{\dagger}(K) \in (1,2)$. Moreover, $\frac{m^*(K)}{K} \to 0$ as $K \to \infty$.*

Proof. Since

$$\lim_{K\to\infty} V^{K}(0) = \frac{(1-\delta)}{1-\delta(1-\lambda)},$$

it follows that

$$\lim_{K\to\infty} m^{\dagger}(K) = \ln\left[\frac{g(1-\delta(1-\lambda))}{1-\lambda}\right] / \ln\delta. \tag{A.5}$$

Therefore, $\frac{m^*(K)}{K} \to 0$ as $K \to \infty$. ■

Proof of Proposition 3

We begin by writing out the value functions of the borrower under the random information structure. Recall that we have $\bar{K}$ memory, and that the borrower is excluded with probability x for $\bar{K}-1$ periods and with complementary probability for $\bar{K}$ periods. The borrower's value function at a clean history has the same form as before:

$$V^{\bar{K},x}(0) = (1-\delta) + \delta\left[\lambda V^{\bar{K},x}(\bar{K}) + (1-\lambda)V^{\bar{K},x}(0)\right].$$

Her value function in the last period of potential exclusion is modified, and is given by:

$$V^{\bar{K},x}(\bar{K}) = \delta^{\bar{K}-1}\ (x\delta + 1 - x)\ V^{\bar{K},x}(0).$$

Using this, we may rewrite her value function at a clean history as

$$V^{\bar{K},x}(0) = \frac{1-\delta}{1-\delta(1-\lambda+\lambda\,\delta^{\bar{K}-1}(x\delta+1-x))}. \tag{A.6}$$

The borrower agrees to repay at G if and only if

$$(1-\delta)g \leq \delta(1-\lambda)\left(1-\delta^{\bar{K}-1}(x\delta+1-x)\right)\ V^{\bar{K},x}(0). \tag{A.7}$$

Using the expression in (A.6), the derivative of the right hand side of (A.7) is

$$\frac{(1-\delta)^3\,(1-\lambda)\,\delta^{\bar{K}}}{\left(1-\delta(1-\lambda)-\lambda\,\delta^{\bar{K}}(x\delta+1-x)\right)^2},$$

which is strictly positive for all $x \in [0,1]$. Thus the right hand side of (A.7) is a strictly increasing function of x. By the definition of $\bar{K}$, (A.7) is satisfied when $x=1$, and violated when $x=0$. There therefore exists $x^*(\bar{K}) \in (0,1]$ such that (A.7) is satisfied for every $x \in [x^*(\bar{K}),1]$.[38]

It remains to examine whether this modification preserves the incentive of lenders not to lend at B. We now show that, while $m^\dagger(\bar{K},x)$ decreases with x, it remains an element of the interval $[1,2)$, so that $m^*(\bar{K},x)=1$ for every $x \in [x^*(\bar{K}),1]$. Thus, the only effect of this modification is to make lending at B less attractive, since the proportion of agents with $m=1$ in the population of agents with a bad signal has been reduced.

We first show that $m^\dagger(\bar{K},x)$ is a strictly decreasing function of x. For every $x \in [x^*(\bar{K}),1]$, $m^\dagger(\bar{K},x)$ is the unique value of $m \in (0,K)$ setting

$$(1-\lambda)(1-x(1-\delta))\left(\delta^{m-1}-\delta^{\bar{K}}\right)V^{\bar{K},x}(0)-(1-\delta)g \tag{A.8}$$

equal to zero. The above is a strictly decreasing function of m. It is also a strictly decreasing function of x, as its derivative with respect to x equals

$$(1-\lambda)\left(\delta^{m-1}-\delta^{\bar{K}}\right)\left[-(1-\delta)\ V^{\bar{K},x}(0)+(1-x(1-\delta))\ \frac{\partial}{\partial x}V^{\bar{K},x}(0)\right]<0,$$

[38] In the non-generic case where $\bar{K}$ satisfies (1) with equality, we have that $x^*(\bar{K})=1$.

where the inequality follows from the fact that $V^{\bar{K},x}(0)$ is positive and strictly decreasing in x. To maintain the expression in (A.8) equal to zero, any increase in x must be compensated by a decrease in $m^\dagger$. The result follows.

We now show that $m^\dagger(\bar{K}, x) \in [1, 2)$ for every $x \in [x^*(\bar{K}), 1]$. Since (A.8) is strictly decreasing in m, it suffices to show that

$$(1-\lambda)(1-x^*(1-\delta))\left(\delta^{2-1}-\delta^{\bar{K}}\right)V^{\bar{K},x^*}(0) < (1-\delta)g, \tag{A.9}$$

and that

$$(1-\lambda)(1-x^*(1-\delta))\left(\delta^{1-1}-\delta^{\bar{K}}\right)V^{\bar{K},x^*}(0) \geq (1-\delta)g. \tag{A.10}$$

Since (A.8) is strictly decreasing in x, we have

$$(1-\lambda)(1-x^*(1-\delta))\left(\delta-\delta^{\bar{K}}\right)V^{\bar{K},x^*}(0) < \delta(1-\lambda)\left(1-\delta^{\bar{K}-1}\right)V^{\bar{K},0}(0),$$

and since the right hand side of (A.7) is strictly increasing in x,

$$\delta(1-\lambda)\left(1-\delta^{\bar{K}-1}\right)V^{\bar{K},0}(0) < \delta(1-\lambda)\left(1-\delta^{\bar{K}-1}(x^*\delta+1-x^*)\right)V^{\bar{K},x^*}(0).$$

By the definition of $x^*(\bar{K})$, the right hand side above equals $(1-\delta)g$, establishing (A.9).

Similarly, since (A.8) is strictly decreasing in x, we have

$$(1-\lambda)(1-x^*(1-\delta))\left(1-\delta^{\bar{K}}\right)V^{\bar{K},x^*}(0) > \delta(1-\lambda)\left(1-\delta^{\bar{K}}\right)V^{\bar{K},1}(0),$$

where $V^{\bar{K},1}(0) = V^{\bar{K}}(0)$, so that the right hand side above is weakly greater than $(1-\delta)g$ by the definition of $\bar{K}$. This establishes (A.10), and completes the proof.

A.5 Non-stationary Analysis

Consider the simple partition of Section 4.1 and the pure strategy equilibrium set out there. Suppose that the game starts at date $t = 1$ with all borrowers having a clean history. The measure of type m borrowers at date t, $\tilde{\mu}_m(t)$, varies over time. Types $m \leq m^*(K)$ repay when extended a loan. Thus, at any date t, the probability that a loan extended at history B is repaid is

$$\tilde{\pi}(t;K) = \frac{\sum_{m=1}^{m^*(K)}\tilde{\mu}_m(t)}{1-\tilde{\mu}_0(t)}. \tag{A.11}$$

The sufficient incentive constraint ensuring that a lender should never want to lend to a borrower with signal B at any date, is

$$\bar{\pi}(K) := \sup_t \tilde{\pi}(t; K) < \frac{\ell}{1+\ell}. \tag{A.12}$$

To compute $\bar{\pi}(K)$, observe that the measure of type $m = 0$ is maximal at $t = 1$. A fraction λ of those borrowers transit to $m = K$, and then transit deterministically thought the lower values of m. Therefore, the measure of type $m = 1$ is maximal at date $K+1$, and equals λ. This is also the date where the repayment probability is maximal. At that date,

$$\tilde{\mu}_m(t) = \begin{cases} \lambda(1-\lambda)^{m-1} & \text{for } m \geq 1, \\ (1-\lambda)^K & \text{for } m = 0, \end{cases}$$

so that

$$\bar{\pi}(K) = \frac{\sum_{m=1}^{m^*(K)} \lambda(1-\lambda)^{m-1}}{1-(1-\lambda)^K} = \frac{1-(1-\lambda)^{m^*(K)}}{1-(1-\lambda)^K}.$$

Letting $m^\infty := \lim_{K\to\infty} m^*(K)$, we have that as $K \to \infty$ the maximal repayment probability converges to

$$1-(1-\lambda)^{m^\infty+1}.$$

Consider the case of $\bar{K}$-period memory, where $m^*(\bar{K}) = 1$. In this case,

$$\bar{\pi}(\bar{K}) = \frac{\lambda}{1-(1-\lambda)^{\bar{K}}}.$$

Note that $\bar{\pi}(\bar{K}) > \frac{1}{\bar{K}}$ (the steady state repayment probability) and that $\bar{\pi}(\bar{K}) \to \frac{1}{\bar{K}}$ as $\lambda \to 0$.

Consequently, if $\bar{\pi}(\bar{K}) < \frac{\ell}{1+\ell}$, then we have a pure strategy equilibrium where the lender's do not have an incentive to lend to a borrower with signal B at any date. If this condition is violated, but the lender's incentive constraint is satisfied in the steady state, then the transition to the steady state is more complex, and requires mixed strategies along the path, and we do not investigate this here.

A.6 Proof of Claims in Section 5.2

We now prove claims related to the mixed strategy equilibrium.

A.6.1 First Claim

We show that $\pi(p, 1)$, defined in (15), is a strictly decreasing function of p. (Continuity for $p \in [0, 1]$ is immediate.) For any given $K \geq 2$ and $p \in (0, 1)$, the invariant distribution $\mu(p, 1)$, is given by equations (12) to (14), together with the condition $\sum_{m=0}^{K} \mu_m = 1$. Solving the above system, we obtain

$$\pi(p, 1) = \frac{p\,(1-p)^{K-1}}{1-(1-p)^K},$$

so that

$$\frac{\partial \pi(p, 1)}{\partial p} = \frac{(1-p)^{K-2}\; h(p, K)}{\left(1-(1-p)^K\right)^2},$$

where $h(p, K) := 1 - Kp - (1-p)^K$ satisfies, for every $p \in (0, 1)$ and $K \geq 1$,

$$h(p, K+1) - h(p, K) = -p\left(1-(1-p)^K\right) < 0,$$

while for every $p \in (0, 1)$, $h(p, 1) = 0$.

We therefore have that for every $p \in (0, 1]$ and $K \geq 2$, $h(p, K) < 0$ so that $\frac{\partial \pi(p,1)}{\partial p} < 0$ and $\pi(p, 1)$ is a strictly decreasing function of p.

A.6.2 Second Claim

We now establish the second equality in (16). For any given $K \geq 2$ and $(p, q) \in (0, 1)^2$, the invariant distribution $\mu(p, q)$, is given by

$$\mu_0 = q(1-\lambda)\,\mu_0 + \Big(1 - p\,(1 - q(1-\lambda))\Big)\,\mu_1,$$

$$\mu_m = (1-p)^{K-m}\mu_K, \quad 1 \leq m \leq K,$$

$$\mu_K = \big(1 - q(1-\lambda)\big)\,\mu_0 + p\,(1 - q(1-\lambda))\;\mu_1 + p\;\sum_{m=2}^{K} \mu_m,$$

together with the condition $\sum_{m=0}^{K} \mu_m = 1$. Solving the above system, we obtain

$$\pi(\tilde{p}, q) = q\,\frac{\tilde{p}\,(1-\tilde{p})^{K-1}}{1-(1-\tilde{p})^K} = q\,\pi(\tilde{p}, 1),$$

as in (16).

A.6.3 Purification of the Mixed Equilibrium in Proposition 4

In the perturbed version of the trust stage game, without loss of generality, it suffices to perturb the payoff to one of the two actions of each of the players. Accordingly, we assume that the payoff to the lender from lending is augmented by εy, where y is the realization of a random variable that is distributed on a bounded support, say $[0, 1]$ (without loss of generality) with a continuous cumulative distribution function, F_Y. The expected payoff to the borrower from wilful default is augmented by εz, where z is the realization of a random variable that is distributed on $[0, 1]$ with a continuous cumulative distribution function, F_Z.

The proof of Lemma 2 is straightforward. Let σ be a stationary sequentially strict equilibrium where each player from population 1 plays the same strategy, and each player from population 2 plays the same strategy. At any information set, since a player has strict best responses, if ε is small enough, then this best response is also optimal for all realizations of the player's payoff shock. Since memory is bounded, there are finitely many strategically distinct information sets for each player. Thus there exists $\bar{\varepsilon} > 0$, such that if $\varepsilon < \bar{\varepsilon}$, there is an equilibrium in the perturbed game that induces the same behavior as σ.

We now turn to the mixed equilibria of Proposition 4. The case where only the lender mixes and the borrower has strict best responses is more straightforward. So we consider first the case where both lender and borrower mix. The equilibrium is not *regular*, since it is contained in the relative interior of a one-dimensional manifold of equilibria. Thus we cannot directly invoke, for example, Doraszelski and Escobar (2010), who show that regular Markov perfect equilibria are purifiable in stochastic games.

Assume that $\ell \in (0, \ell^*)$, so that in the unperturbed game, the mixed equilibrium has the lender lending with probability $\tilde{p}$ after credit history B, while the borrower repays with probability $\tilde{q}$ if $m \in \{0, 1\}$ and has strict incentives to default if $m > 1$. Now if ε is small enough, and if the lender's lending probability after B is close to $\tilde{p}$, the borrower retains strict incentives to default when $m > 1$ for every realization z. Similarly, the lender retains strict incentives to lend after signal G. Let $\bar{y}$ denote the threshold value of the payoff shock, such that the lender lends after signal B if and only if $y > \bar{y}$, and let $\bar{p} := 1 - F_Y(\bar{y})$. Let $\bar{z}$ denote the threshold value of the payoff shock, such that a borrower with $m \in \{0, 1\}$ defaults if and only if $z > \bar{z}$, and define $\bar{q} := F_Z(\bar{z})$. At $(p, q) = (\bar{p}, \bar{q})$, the value functions in the perturbed game can then be rewritten to take into account the payoff shocks. For $m = 0$, we have

$$\tilde{V}^{\bar{K}}(0, \bar{p}) = (1 - \delta)\left(1 + \varepsilon \int_{\bar{z}}^{1} (z - \bar{z})\, dF_Z(z)\right) + \delta \left[\lambda \tilde{V}^{\bar{K}}(\bar{K}, \bar{p}) + (1 - \lambda)\tilde{V}^{\bar{K}}(0, \bar{p})\right]. \quad \text{(A.13)}$$

We derive the above expression as follows. At $\bar{z}$, the borrower is indifferent between defaulting and repaying. Thus the payoff from defaulting, when $z > \bar{z}$, equals the payoff from repaying plus the difference $z - \bar{z}$. Similarly, the value function of a borrower with signal B and $m = 1$ is given by

$$\tilde{V}^{\bar{K}}(1, \bar{p}) = \bar{p}(1-\delta)\left(1+\varepsilon \int_{\bar{z}}^{1} (z-\bar{z})\, dF_Z(z)\right) + \delta \left[\bar{p}\lambda \tilde{V}^{\bar{K}}(\bar{K}, \bar{p}) + (1-\bar{p}\lambda)\tilde{V}^{\bar{K}}(0, \bar{p})\right]. \quad \text{(A.14)}$$

For $m > 1$, the borrower always defaults, and hence

$$\tilde{V}^{\bar{K}}(m, \bar{p}) = \bar{p}(1-\delta)\left(1 + g + \varepsilon\, \mathbb{E}(z)\right) + \delta \left[\bar{p}\tilde{V}^{\bar{K}}(\bar{K}, \bar{p}) + (1-\bar{p})\tilde{V}^{\bar{K}}(m-1, \bar{p})\right] \quad \text{(A.15)}$$

The indifference condition for a borrower with $m \in \{0, 1\}$ and $z = \bar{z}$ is

$$(1-\delta)(g + \varepsilon \bar{z}) - \delta(1-\lambda)\left(\tilde{V}^{\bar{K}}(0, \bar{p}) - \tilde{V}^{\bar{K}}(\bar{K}, \bar{p})\right) = 0. \quad \text{(A.16)}$$

The indifference condition for a lender facing a borrower with history B and having payoff shock $\bar{y}$ is

$$\frac{\bar{q}\mu_1(\bar{p}, \bar{q})}{1 - \mu_0(\bar{p}, \bar{q})} - \frac{\ell}{1 + \ell + \varepsilon \bar{y}} = 0. \quad \text{(A.17)}$$

When $\varepsilon = 0$, equations (A.16) and (A.17) have as solution $(\tilde{p}, \tilde{q})$. In the remainder of this appendix, we establish that the Jacobian determinant of the left hand side of equations (A.16) and (A.17) at $\varepsilon = 0$ and $(p, q) = (\tilde{p}, \tilde{q})$ is non-zero. By the implicit function theorem, if ε is sufficiently close to zero, there exist $(\bar{p}(\varepsilon), \bar{q}(\varepsilon))$ close to $(\tilde{p}, \tilde{q})$ that solves equations (A.16) and (A.17). We have then established that the mixed strategy equilibrium where both the lender and the borrower mix is purifiable.

The indifference condition for a borrower with type $m \in \{0, 1\}$ in the unperturbed game is given by[39]

$$\gamma(p, q) = (1-\delta)g - \delta(1-\lambda)(V^{\bar{K}}(0, p, q) - V^{\bar{K}}(\bar{K}, p, q)) = 0. \quad \text{(A.18)}$$

The indifference condition for the lender after B in the unperturbed game is given by

$$\phi(p, q) = q\mu_1(p, q)(1 + \ell) - \ell(1 - \mu_0(p, q)) = 0. \quad \text{(A.19)}$$

[39] Although, for every $m \in \{0, 1, \ldots, \bar{K}\}$, the level of $V^{\bar{K}}(m, p)$ is independent of q when $p = \tilde{p}$, its slope is not. We therefore emphasize the dependence of $V^{\bar{K}}$ on q in the remainder of this section.

Consider the partial derivatives, $\phi_p, \phi_q, \gamma_p, \gamma_q$, as 2×2 matrix. We now prove that the determinant of this matrix is non-zero when evaluated at $(p,q) = (\tilde{p}, \tilde{q})$. The value functions of the borrower evaluated at $p = \tilde{p}$ are constant with respect to q. Therefore, $\gamma_q = 0$ when $p = \tilde{p}$. Thus it suffices to prove that γ_p and ϕ_q are both non-zero at $(p,q) = (\tilde{p}, \tilde{q})$.

When $q \in (0,1)$, $V^{\bar{K}}(0,p,q)$ and $V^{\bar{K}}(K,p,q)$ satisfy

$$V^{\bar{K}}(0,p,q) = (1-\delta)(1+g(1-q)) + \delta\Big((1-(1-\lambda)q)\ V^{\bar{K}}(K,p,q) + (1-\lambda)q\ V^{\bar{K}}(0,p,q)\Big).$$

Differentiating with respect to p, we obtain

$$\frac{\partial V^{\bar{K}}(K,p,q)}{\partial p} = \frac{1-\delta(1-\lambda)q}{\delta - \delta(1-\lambda)q}\ \frac{\partial V^{\bar{K}}(0,p,q)}{\partial p},$$

where $(1-\delta(1-\lambda)q)/(\delta - \delta(1-\lambda)q) > 1$. As a result,

$$\gamma_p(p,q) = -\delta(1-\lambda)\left[1 - \frac{1-\delta(1-\lambda)q}{\delta(1-(1-\lambda)q)}\right] \frac{\partial V^{\bar{K}}(0,p,q)}{\partial p}$$

is strictly positive, since $V^{\bar{K}}(0,p,q)$ is a strictly increasing function of p for every $q \in [0,1]$. Thus γ_p is non-zero at $(p,q) = (\tilde{p}, \tilde{q})$.

Differentiating ϕ with respect to q gives

$$\phi_q(p,q) = \frac{\partial \mu_1(p,q)}{\partial q} q(1+\ell) + \mu_1(p,q)(1+\ell) + \frac{\partial \mu_0(p,q)}{\partial q}\ell. \tag{A.20}$$

Solving the system for the invariant distribution of types, we have

$$\mu_0(p,q) = \frac{C(1-p+pQ)}{A-QB},$$

$$\mu_1(p,q) = \frac{C(1-Q)}{A-QB},$$

where

$$A = (2-p)C + S, \quad B = (1-p)C + S, \quad C = (1-pS),$$

$$Q = q(1-\lambda), \quad S = \sum_{m=2}^{k} (1-p)^{k-m}.$$

Differentiating with respect to q,

$$\frac{\partial \mu_0(p,q)}{\partial q} = \frac{(1-\lambda)C}{(1-Q)(A-QB)}\ (1-\mu_0(p,q))\,,$$

$$\frac{\partial \mu_1(p,q)}{\partial q} = \frac{-(1-\lambda)C}{(1-Q)(A-QB)}\ \mu_1(p,q).$$

Using these expressions in (A.20), we obtain

$$\phi_q(p,q) = \mu_1(p,q)(1+\ell)\left[1 - \frac{QC}{(1-Q)(A-QB)}\right] + \frac{(1-\lambda)C}{(1-Q)(A-QB)}\ (1-\mu_0(p,q))\,\ell. \tag{A.21}$$

Since the lender is indifferent between lending and not lending at a bad history when $q = \tilde{q}$, we have that for every $p \in [0,1]$,

$$\mu_1(p,\tilde{q})(1+\ell)\tilde{q} = (1-\mu_0(p,\tilde{q}))\,\ell.$$

Using this indifference condition in (A.21) gives

$$\phi_q(p,q)\big|_{q=\tilde{q}} = \mu_1(p,\tilde{q})(1+\ell) > 0$$

for every $p \in [0,1]$, and we have established that ϕ_q is non-zero at $(p,q) = (\tilde{p},\tilde{q})$.

Finally, the equilibrium where only the lender mixes and the borrower has strict best responses is also purifiable, since we have shown that $\gamma_p(p,1) \neq 0$.

A.7 General Games

Let Γ be a two-player (stage) game of perfect information, with finitely many nodes and no chance moves. Let Z be the set of terminal nodes or outcomes, so that each element $z \in Z$ is associated with a utility pair, $u(z) \in \mathbb{R}^2$. Assume that there are no payoff ties, so that if $z \neq z'$, then $u_i(z) \neq u_i(z')$ for every $i \in \{1,2\}$. Thus there exists a unique backwards induction strategy profile, $\bar{\sigma}$ and a unique backwards induction outcome, denoted $\bar{z}$. Normalize payoffs so that $u(\bar{z}) = (0,0)$ and $u(z^*) = (1,1)$.

Let X denote the set of non-terminal nodes, partitioned into X_1 and X_2, the decision nodes of the two players. Any pure behavior strategy profile σ induces a terminal node starting at any non-terminal node x. Write $u(\sigma(x))$ for the payoffs so induced. For any $x \in X$, let $\bar{\sigma}(x)$ denote the unique backwards induction path induced by $\bar{\sigma}$ starting at x, and

$u(\bar{\sigma}(x))$ denote the payoff vector at the corresponding terminal node. Given any terminal node $z \in Z$, let $\phi(z)$ denote the path to z from the initial node x_0.

Definition 4 *Fix a terminal node z, a path $\phi(z)$, and a node x on this path, where player i moves. Player i has an incentive to deviate at x if $u_i(\bar{\sigma}(x)) > u_i(z)$. Player i has an incentive to deviate from $\phi(z)$ if there exists a node x on this path where he has an incentive to deviate.*

Remark 5 *No player has an incentive to deviate from $\phi(\bar{z})$, the backwards induction path. If $z \neq \bar{z}$, then at least one player has an incentive to deviate from the path to z.*

We focus on the sustainability of outcomes that Pareto-dominate the backwards induction outcome $\bar{z}$. Consider any pair (Γ, z^*), where Γ is a generic two-player game and z^* is a terminal node that strictly Pareto-dominates the backwards induction outcome $\bar{z}$, i.e. where $u_1(z^*) > 0$ and $u_2(z^*) > 0$. We assume that the pair (Γ, z^*) satisfies the following assumptions:

- Only one player (labelled $\hat{i}$) has an incentive to deviate on the path to z^*.
- There is a single node $\hat{x}$ on $\phi(z^*)$ at which $\hat{i}$ has an incentive to deviate.

Let the pair (Γ, z^*) satisfy the two assumptions, and let j index the player who does not have an incentive to deviate. If $\hat{i}$ initiates the play of the backwards induction profile at $\hat{x}$, and j continues, the resulting payoff $u_{\hat{i}}(\bar{\sigma}(\hat{x}))$ to $\hat{i}$ is strictly greater than her payoff $u_{\hat{i}}(z^*)$ at z^*, which we normalizes to 1. Thus we may write $1 + g$ for this payoff, where

$$g := u_{\hat{i}}(\bar{\sigma}(\hat{x})) - u_{\hat{i}}(z^*) > 0.$$

Define $\tilde{x}_j$ as the maximal element under the precedence relation $\preceq$ on X (where for $x, y \in X$, $x \preceq y$ signifies that x precedes y) of the set $\tilde{X}_j$, where

$$\tilde{X}_j = \{ x \in X_j \cap \phi(z^*), \, x \preceq \hat{x}, \, u_j(\bar{\sigma}(x)) > u_j(\bar{\sigma}(\hat{x})) \} .$$

For instance, in the centipede game depicted in Figure 9, we have $\hat{x} = (r_1, r_2, r_3)$ and $\tilde{x}_j = (r_1, r_2)$.

We show that the set $\tilde{X}_j$ is non-empty, so that $\tilde{x}_j$ is well defined. If $\tilde{X}_j$ is empty, this implies that, for all $x \in X_j \cap \phi(z^*)$ such that $x \preceq \hat{x}$, we have $u_j(\bar{\sigma}(x)) \leq u_j(\bar{\sigma}(\hat{x}))$. Since player $\hat{i}$'s incentive to deviate from $\phi(z^*)$ is maximal at $\hat{x}$, we have that $u_{\hat{i}}(\bar{\sigma}(x)) \leq u_{\hat{i}}(\bar{\sigma}(\hat{x}))$ for every $x \in X_{\hat{i}} \cap \phi(z^*)$ such that $x \preceq \hat{x}$. These two facts imply that $\bar{\sigma}(\hat{x})$ is the backwards induction outcome, $\bar{z}$. But since, by the definition of $\hat{x}$, $u_{\hat{i}}(\bar{\sigma}(\hat{x})) > u_{\hat{i}}(z^*)$, this contradicts the assumption that z^* Pareto-dominates $\bar{z}$.

We may therefore define

$$\ell := u_j(\bar{\sigma}(\tilde{x}_j)) - u_j(\bar{\sigma}(\hat{x})) > 0.$$

In words, ℓ is player j's loss from continuing on the path $\phi(z^*)$ at $\tilde{x}_j$ if player $\hat{i}$ continues with her backwards induction strategy at $\hat{x}$. The argument above established that $\ell > 0$.

The information structure in the repeated random matching game is a generalization of the simple information structure that has been extensively used in this paper in the context of the borrower-lender game. Partition the set of terminal nodes Z in the stage game so that $\mathcal{D}$ denotes the set of nodes that arises after a deviation by $\hat{i}$ from $\phi(z^*)$ at $\hat{x}$. Let $\mathcal{N}$ denote the complement, $\mathcal{N} = Z \setminus \mathcal{D}$. The signal regarding player $\hat{i}$ is B if there is any instance of $\mathcal{D}$ in any of the last K periods; otherwise, the signal is G. In each period, player j observes the signal regarding $\hat{i}$ before the players play the stage game Γ. Player $\hat{i}$ observes no information regarding the past play of player j. Given this information structure, let $m \in \{1, \ldots, K\}$ denote the number of periods that must elapse before player $\hat{i}$'s signal switches back to G given that it is currently B. Under this information structure, m is player $\hat{i}$'s private information, i.e. her type. Let $m^* \in (1, K)$ be a threshold used in defining player $\hat{i}$'s strategy. In equilibrium, it will depend on player $\hat{i}$'s payoff function and will generically take non-integer values.

Define the following strategies and strategy profiles in the stage game Γ. Let $\sigma^* = (\sigma^*_{\hat{i}}, \sigma^*_j)$ denote the strategy in Γ where $\phi(z^*)$ is played unless some player deviates from $\phi(z^*)$, in which case players continue with $\bar{\sigma}$.

For $i \in \{1, 2\}$, define the strategy $\hat{\sigma}_i$ in Γ as follows. For every $x \in X_i$,

$$\hat{\sigma}_i(x) = \begin{cases} \sigma^*_i(x) & \text{if } x \in \phi(z^*) \text{ and } x \succeq \hat{x}, \\ \bar{\sigma}_i(x) & \text{otherwise.} \end{cases}$$

The repeated game strategies are as follows.

- The players play σ^* at G.
- Player j plays $\hat{\sigma}_j$ at B.
- Player $\hat{i}$ plays $\bar{\sigma}_{\hat{i}}$ at B if $m > m^*$ and plays $\hat{\sigma}_{\hat{i}}$ at B if $m < m^*$.

Given these strategies, it is straightforward to verify that the value of player $\hat{i}$ at signal G is $V^K(0) := 1$, while her value at signal B, given her type m, is $V^K(m) := \delta^m$. Since the deviation gain for $\hat{i}$ equals g, then if player $\hat{i}$'s discount factor δ is large enough, there exists $\bar{K}$ such that if $K \geq \bar{K}$, player $\hat{i}$ has no incentive to deviate from $\phi(z^*)$ when she has a good signal.

We now verify the optimality of these repeated game strategies. Consider signal G. If $K \geq \bar{K}$, then player $\hat{i}$ has no incentive to deviate from σ^* at G. Given this, neither does player j, since by assumption, j does not have an incentive to deviate from $\phi(z^*)$.

Now consider signal B, and a type m for player $\hat{i}$. Consider any node $x \preceq \tilde{x}_j$ on the path $\phi(z^*)$. Given that player j plays $\bar{\sigma}_j$ at $\tilde{x}_j$, backwards induction establishes that $\bar{\sigma}_i$ is optimal for $i \in \{1, 2\}$ at every node x that precedes $\hat{x}$.

Consider next the node $\hat{x}$. If player $\hat{i}$ plays $\bar{\sigma}_{\hat{i}}$ at this node, the specified continuation strategies imply that she gets a current payoff of $u_{\hat{i}}(\bar{\sigma}(\hat{x}))$ and a continuation value of $V^K(K)$. If instead she continues on path $\phi(z^*)$, she gets a current payoff of $u_{\hat{i}}(z^*)$ and a continuation value of $V^K(m-1) = \delta^{m-1}$. Thus the payoff difference between these two choices equals

$$(1-\delta)g - \delta[V^K(m-1) - V^K(K)].$$

Since $V^K(m)$ is strictly decreasing in m, there exists a real number m^* such that at $\hat{x}$ it is optimal for $\hat{i}$ to continue with $\sigma^*_{\hat{i}}$ if $m < m^*$ and with $\bar{\sigma}_{\hat{i}}$ otherwise. Furthermore, since the deviation gain for $\hat{i}$ is maximal at $\hat{x}$, it is optimal to also continue with $\sigma^*_{\hat{i}}$ at subsequent nodes on the path $\phi(z^*)$ if $m < m^*$.

There remains the critical node $\tilde{x}_j$. If j continues on the path $\phi(z^*)$ at this node, and $m > m^*$ so that $\hat{i}$ proceeds with her backwards induction strategy, j incurs a strict loss relative to playing his backwards induction strategy at $\tilde{x}_j$, since $u_j(\bar{\sigma}(\hat{x})) - u_j(\bar{\sigma}(\tilde{x}_j)) = -\ell$. On the other hand, if $m < m^*$ so that $\hat{i}$ continues on the path $\phi(z^*)$ at $\hat{x}$, j continuing on $\phi(z^*)$ at $\tilde{x}_j$ rather than playing the backwards induction strategy secures the net gain $u_j(z^*) - u_j(\bar{\sigma}(\tilde{x}_j)) > 0$. (This payoff difference is positive because we assumed that j has no incentive to deviate from $\phi(z^*)$.)

Let π denote the probability assigned by j to player $\hat{i}$'s type m being strictly less than m^*. Then it is optimal for j to play his backwards induction strategy at node $\tilde{x}_j$ if

$$\pi < \frac{u_j(\bar{\sigma}(\tilde{x}_j)) - u_j(\bar{\sigma}(\hat{x}))}{u_j(z^*) - u_j(\bar{\sigma}(\hat{x}))} = \frac{\ell}{1 - u_j(\bar{\sigma}(\tilde{x}_j)) + \ell} =: \bar{\pi}.$$

Suppose that the pair (Γ, z^*) also satisfies the third assumption set out in Section 7.3, i.e.

- Player $\hat{i}$ has a decision node x that precedes $\hat{x}$ on the path to z^* where her backwards induction strategy prescribes deviating from $\phi(z^*)$ at x.

In this case it is costly for $\hat{i}$ to continue on the path $\phi(z^*)$ at x, given that j plays the backwards induction strategy at $\tilde{x}_j$. Thus the D1 criterion implies that if the decision node $\tilde{x}_j$ is reached, the probability π assigned by j must be zero. The arguments for this are identical to those set out in the proof of Proposition 6. Thus we have an equilibrium that supports the outcome z^* without any further assumptions.

If the above assumption is not satisfied, then π is determined by the invariant distribution over the values of m. Ensuring that π is low enough requires arguments similar to those explored in the context of the basic lender-borrower game where the borrower need not make a prior application, and we do not repeat them here.

A.8 An Alternative Specification of Forgetting

We now consider an alternative modeling of forgetting defaults and show that it yields qualitatively similar conclusions. Suppose that the borrower's history is edited, so that any incidence of $\mathcal{D}$ is replaced by $\mathcal{N}$ after K periods have elapsed, i.e. it is as though the loan never took place. Let us also assume that in each period there is a small probability ρ that a borrower does not meets a lender. When a lender and a borrower are matched, the lender perfectly observes the borrower's entire *edited* history, and that $K \geq \bar{K}$, so that an incidence of default is retained longer in the record than is required for incentivizing the borrower.

Let $\tilde{H}^t = \mathcal{O}^t$ denote the space of t-period private histories of the borrower. These are known only to the borrower. Let H^t denote the space of t-period recorded histories. These are the edited histories observed by the lender. If $t \leq K$, then $H^t = \mathcal{O}^t$, while it $t > K$, then $H^t = \{\mathcal{R}, \mathcal{N}\}^{t-K} \times \mathcal{O}^K$. The following lemma shows that in any purifiable equilibrium, the borrower will not condition her strategy on her private history, but only on the recorded history. Indeed, the former is payoff irrelevant since no lender that she will ever be matched with has access to it. Consequently, whenever we speak of the history in the rest of this subsection, we mean the recorded history. We now show that there are further restrictions on how the recorded history may be utilized in any purifiable equilibrium.

We define the following equivalence relation on t-period histories. Consider two histories, $h^t = (a_1, ...a_t)$ and $\hat{h}^t = (\hat{a}_1, ..., \hat{a}_t)$. We write that $h^t \sim \hat{h}^t$ if for every $s \in \{1, \ldots, t-K\}$,

$\hat{a}_s = a_s$, while for every $s \in \{t-K+1, \ldots, t\}$, $\hat{a}_s \neq a_s \Rightarrow \hat{a}_s, a_s \in \{\mathcal{D}, \mathcal{N}\}$. That is, two t-period histories are equivalent if:

- The outcomes are identical in any period $s \leq t - K$.
- If the outcomes differ in any period s within the last K periods, then these outcomes are either $\mathcal{D}$ or $\mathcal{N}$.

Let σ denote a strategy profile in Γ^∞, i.e. a strategy for borrowers and a strategy for lenders.

Lemma A.2 *If σ is a purifiable equilibrium, then at every date $t+1$, σ is measurable with respect to H^t, the set of possible recorded histories. Further, if $h^t \sim \hat{h}^t$, $\sigma(h^t) = \sigma(\hat{h}^t)$.*

Proof. The strategy of any lender that the borrower meets at any future date cannot condition on the private history $\tilde{h}^t$. Thus the borrower's continuation value does not depend $\tilde{h}^t$, and nor does his current payoff. Hence the borrower can only condition upon $\tilde{h}^t$ if he is indifferent between R and D. However, in the perturbed game, such indifference is possible only for a set of z values that has Lebsegue measure zero. Thus any equilibrium in the unperturbed game where the borrower chooses different mixed actions after different private histories is not purifiable.

The proof of the second part of the lemma is by induction. Let h^t, $\hat{h}^t$ be two recorded histories that are equivalent. At all dates $s > t+K$ the lenders will not be able to condition upon these histories since they will not be distinguishable. Thus in period $s-1$, the borrower will not condition her repayment decision on these histories. As a result, the lender at date $s-1$ will also not condition her lending decision on these histories. By induction, neither lender nor borrower at any previous date will condition their behaviors on these histories. ∎

This lemma implies that in any purifiable equilibrium, not obtaining a loan must be treated in the same way as a default. This has important implications. Consider an equilibrium where borrowers are incentivized to repay by $\bar{K}$ periods of exclusion. The above lemma implies that if any period, a borrower fails to get a loan, then she must also be excluded for $\bar{K}$ periods. Thus if $\rho > 0$, so that there is some chance that a borrower might fail to get a loan for exogenous reasons, the fact that such failures must lead to punishment causes additional inefficiencies. Providing the lender with coarse information can improve efficiency in this context as well. In particular, if the borrower is only informed about the last $\bar{K}$

outcomes, and only learns whether the borrower has defaulted or not in this time, then an equilibrium can be sustained under exactly the same conditions as in Section 4.1.

References

Ali, S. N., and D. A. Miller (2013): "Enforcing cooperation in networked societies," *Unpublished paper.[282].*

Bergemann, D., and S. Morris (2016): "Bayes correlated equilibrium and the comparison of information structures in games," *Theoretical Economics*, 11(2), 487–522.

Bhaskar, V. (1998): "Informational Constraints and the Overlapping Generations Model: Folk and Anti-Folk Theorems," *Review of Economic Studies*, 65(1), 135–149.

Bhaskar, V., G. J. Mailath, and S. Morris (2013): "A Foundation for Markov Equilibria in Sequential Games with Finite Social Memory," *Review of Economic Studies*, 80(3), 925–948.

Bulow, J., and K. Rogoff (1989): "A constant recontracting model of sovereign debt," *Journal of Political Economy*, 97(1), 155–178.

Camera, G., and M. Casari (2009): "Cooperation among strangers under the shadow of the future," *The American Economic Review*, 99(3), 979–1005.

Cho, I.-K., and D. Kreps (1987): "Signaling Games and Stable Equilibria," *Quarterly Journal of Economics*, 102(2), 179–221.

Deb, J. (2008): "Cooperation and Community Responsibility: A Folk Theorem for Repeated Random Matching Games," Discussion paper, mimeo.

Deb, J., and J. González-Díaz (2010): "Community enforcement beyond the prisoners dilemma," Discussion paper, Working paper, July.

Dobbie, W., P. Goldsmith-Pinkham, N. Mahoney, and J. Song (2016): "Bad credit, no problem? Credit and labor market consequences of bad credit reports," Discussion paper, National Bureau of Economic Research.

Doraszelski, U., and J. Escobar (2010): "A Theory of Regular Markov Perfect Equilibria in Dynamic Stochastic Games: Genericity, Stability, and Purification," *Theoretical Economics*, 5(3), 369–402.

Duffy, J., and J. Ochs (2009): "Cooperative behavior and the frequency of social interaction," *Games and Economic Behavior*, 66(2), 785–812.

Dutta, S. (1992): "Building Trust," Mimeo, London School of Economics.

Ekmekci, M. (2011): "Sustainable Reputations with Rating Systems," *Journal of Economic Theory*, 146(2), 479–503.

Ellison, G. (1994): "Cooperation in the Prisoner's Dilemma with Anonymous Random Matching," *Review of Economic Studies*, 61(3), 567–588.

Elul, R., and P. Gottardi (2015): "Bankruptcy: Is It Enough to Forgive or Must We Also Forget?," *American Economic Journal: Microeconomics*, 7(4), 294–338.

Ely, J. C. (2017): "Beeps," *The American Economic Review*, 107(1), 31–53.

Ely, J. C., and J. Välimäki (2002): "A Robust Folk Theorem for the Prisoner's Dilemma," *Journal of Economic Theory*, 102(1), 84–105.

Fudenberg, D., D. Kreps, and E. Maskin (1990): "Repeated Games with Long-Run and Short-Run Players," *Review of Economic Studies*, 57(4), 555–574.

Ghosh, P., and D. Ray (1996): "Cooperation in community interaction without information flows," *The Review of Economic Studies*, 63(3), 491–519.

Govindan, S., P. J. Reny, and A. J. Robson (2003): "A Short Proof of Harsanyi's Purification Theorem," *Games and Economic Behavior*, 45(2), 369–374.

Gross, T., M. J. Notowidigdo, and J. Wang (2016): "The marginal propensity to consume over the business cycle," Discussion paper, National Bureau of Economic Research.

Harsanyi, J. C. (1973): "Games with Randomly Disturbed Payoffs: A New Rationale for Mixed-Strategy Equilibrium Points," *International Journal of Game Theory*, 2(1), 1–23.

Heller, Y., and E. Mohlin (2017): "Observations on cooperation," *Review of Economic Studies, forthcoming.*

Kamenica, E., and M. Gentzkow (2011): "Bayesian persuasion," *The American Economic Review*, 101(6), 2590–2615.

Kandori, M. (1992): "Social Norms and Community Enforcement," *The Review of Economic Studies*, 59(1), 63–80.

Kocherlakota, N. R. (1998): "Money Is Memory," *Journal of Economic Theory*, 81(2), 232–251.

Kovbasyuk, S., and G. Spagnolo (2016): "Memory and Markets," *Available at SSRN 2756540*.

Kranton, R. E. (1996): "The formation of cooperative relationships," *The Journal of Law, Economics, and Organization*, 12(1), 214–233.

Kremer, I., Y. Mansour, and M. Perry (2014): "Implementing the Wisdom of the Crowd," *Journal of Political Economy*, 122(5), 988–1012.

Liu, Q., and A. Skrzypacz (2014): "Limited Records and Reputation Bubbles," *Journal of Economic Theory*.

Musto, D. K. (2004): "What happens when information leaves a market? Evidence from postbankruptcy consumers," *The Journal of Business*, 77(4), 725–748.

Nava, F., and M. Piccione (2014): "Efficiency in repeated games with local interaction and uncertain local monitoring," *Theoretical Economics*, 9(1), 279–312.

Piccione, M. (2002): "The Repeated Prisoner's Dilemma with Imperfect Private Monitoring," *Journal of Economic Theory*, 102(1), 70–83.

Sperisen, B. (2016): "Bounded Memory, Reputation, and Impatience," Discussion paper.

Sugaya, T. (2013): "Folk Theorem in Repeated Games with Private Monitoring," unpublished, Stanford Graduate School of Business.

Takahashi, S. (2010): "Community Enforcement When Players Observe Partners' Past Play," *Journal of Economic Theory*, 145(1), 42–62.

Wiseman, T. (2015): "A note on the essentiality of money under limited memory," *Review of Economic Dynamics*, 18(4), 881–893.

Wolitzky, A. (2012): "Cooperation with network monitoring," *Review of Economic Studies*, 80(1), 395–427.