

Community Enforcement of Trust*

V Bhaskar[†] Caroline Thomas[‡]

June 9, 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. To incentivise trustworthiness, defaulters should be punished by temporary exclusion. However, it is profitable to trust defaulters who are on the verge of rehabilitation. With perfect bounded information, defaulter exclusion unravels and trust cannot be sustained, in any purifiable equilibrium. A coarse information structure, that pools recent defaulters with those nearing rehabilitation, endogenously generates adverse selection, sustaining punishments. Equilibria where defaulters are trusted with positive probability improve efficiency, by raising the proportion of likely re-offenders in the pool of defaulters.

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, Dean Corbae, Mehmet Ekmekci, Mark Feldman, John Geanakoplos, Andy Glover, Johannes Hörner, George Mailath, Larry Samuelson, Tom Wiseman, and seminar audiences at Austin, Chicago, Harvard-MIT, NYU, Stanford, Turin, Toronto, Toulouse, Western Ontario, and Yale for helpful comments. We thank the editor (Adam Szeidl) and four anonymous referees for many useful suggestions. 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, which 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 behaviour (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 plausible in any context, in many situations — notably, consumer credit markets — this is legally mandated. 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.¹ Bounded memory also arises under policies used by internet platforms to compute the scores summarising 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 municipalities² have introduced “ban the box” legislation, prohibiting employers from asking job applicants about prior convictions unless those relate directly to the job.³

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.⁴ Our model has a large number of long-lived

¹They argue that limited records may be welfare-improving in the presence of adverse selection — see also Kovbasyuk and Spagnolo (2016). We do not provide a rationale for bounded memory, but only examine its implications.

²See “Pandora’s box” in *The Economist*, August 13th 2016.

³The broader philosophical appeal of the principle that an individual’s transgressions should not be perpetually held against them is embodied in the European Court of Justice’s determination that individuals have the “right to be forgotten”, and may compel search engines to delete past records.

⁴This 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.⁵ 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. With bounded memory, *disciplining lenders* to not lend to recently defaulters 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 behaviour 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 purifiable equilibrium — that is, any equilibrium that is robust to small payoff shocks — borrower exclusion unravels.

The intuition is best illustrated if we focus on equilibria where every player has strict incentives at every information set. Observe that 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, in any equilibrium with strict incentives, she will repay a loan whenever 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 argument extends to all equilibria that are robust to small payoff shocks, since such equilibria are essentially strict.

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

⁵Indeed, in a modern economy, borrowers have access to a variety of sources of financing.

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. Our question is, how can coarse information and the consequent adverse selection be tailored to sustain efficient outcomes?

The simple 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 borrower.

Nonetheless, we show that under the simple 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, this is achieved in a pure strategy equilibrium and the lender's profits are also constrained optimal. If loans are very profitable, 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 results in low profits for lenders, by inducing a large pool of borrowers with bad records.

We also show that if the interaction must be initiated, at a small cost, by the party to be trusted (the borrower), full efficiency is ensured. In conjunction with coarse information, such a modification transforms the interaction between borrowers and lenders into a signalling 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, making lenders suspicious.

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 in any purifiable equilibrium. Section 4 shows that a simple 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, and also ensures high lender payoffs if the borrower has to initiate the interaction. 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. Contagion strategies, where a single defection results in the breakdown of cooperation throughout the population, are used in order to support cooperation.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 analyse the prisoner’s dilemma played in a large population. Takahashi 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 assume that a player observes a random sample of the past actions of her opponent, and 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. They 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

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

⁷Experimental 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.

⁸The 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 her opponent. Finite memory belief-free strategies in our setting are not purifiable.

consider stage games with a sequential structure. In the 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. One-sided moral hazard arises naturally in product markets, where sellers have to be incentivised to provide quality, as in Klein and Leffler (1981). It also arises in Greif (1993), who examines mechanisms for deterring opportunism among the Maghribi traders, and in Tirole (1996), who studies the role of collective reputations that are attached to groups. Deb and González-Díaz (2010) analyse simultaneous-move stage games with one-sided moral hazard that are played in a random matching environment. Karlan, Mobius, Rosenblat, and Szeidl (2009) examine the role of friendship ties in sustaining lending in a network, and characterise the maximum flow of lending between lender and borrower when punishments must be coalition-proof.

Liu and Skrzypacz (2014) combine one-sided moral hazard, bounded memory and incomplete information. Buyers are short-lived and have bounded information about the seller’s past decisions, but do not observe the magnitude of 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, where the seller builds up his reputation before milking it.⁹ 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.

A major methodological difference between our work and much of the recent research on repeated games is our insistence on equilibria where players have strict incentives. We require that equilibria be robust to small payoff shocks, and thus be purifiable, as in Harsanyi (1973). We view purifiability as a mild robustness requirement, and in our context it 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 offers 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

⁹Sperisen (2018) extends this analysis by considering non-stationary equilibria.

that purifiability, in conjunction with bounded memory, results in a total breakdown of cooperative behaviour. In contrast, the present paper shows that, by providing partial information on past histories, one can robustly support efficient outcomes. The information structures and equilibrium strategies that sustain efficiency are extremely simple and also intuitive.

A central feature of our analysis is the role of *endogenous adverse selection*. Although the underlying environment has moral hazard but no adverse selection, an optimal information structure does not fully reveal the borrower's recent history to the lender. Endogenous adverse selection also arises in repeated games with private monitoring (Sekiguchi (1997) and Bhaskar and Obara (2002)), where a player randomises over two pure strategies, so that the private signals observed by the opponent are informative of the player's continuation play. Similarly, in Rahman (2012) both the worker and her monitor randomise in order to incentivise each other, to work and to monitor, respectively.

2 The Model

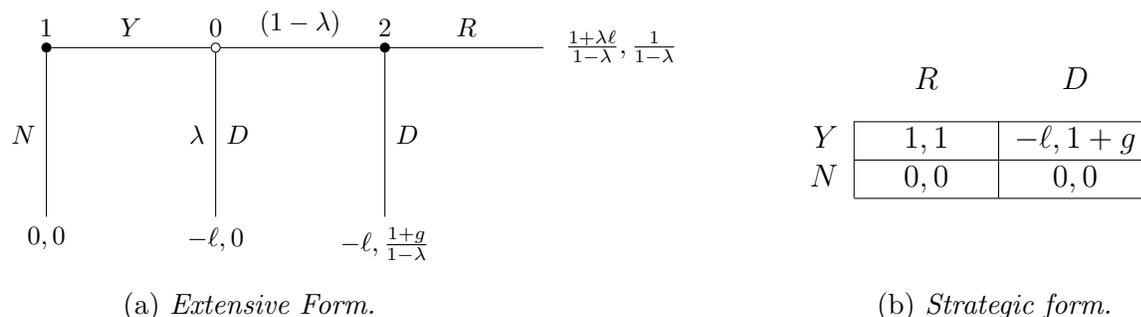


Figure 1: *Extensive and strategic forms of the Trust Game Γ*

Time is discrete and the horizon infinite. In each period, individuals from a continuum population 1 are randomly matched with individuals from a continuum population 2, to play the trust game Γ 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, and is constrained to default. 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 $\ell > 0$, it is a one-sided prisoner’s dilemma. The key features of the trust game are:

- The outcome of the backwards induction profile (N, D) , where player 1 chooses N and player 2 chooses D , is Pareto-dominated by the (random) outcome that results when the players play (Y, R) .
- If players expect (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, realised 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. For concreteness and expositional clarity, we fix on the credit market interpretation.

Let Γ^∞ denote the infinitely repeated game where at every period players are randomly matched to play the trust game Γ . The borrower has a discount factor $\delta \in (0, 1)$. The discount factor of the lender is irrelevant for positive analysis.¹⁰ Since the borrower has a short-term incentive to default, incentives to repay can only be provided by her future lenders. The information that these lenders have about the borrower’s past behaviour will be based on her last K outcomes in \mathcal{O} , where $\mathcal{O} = \{\mathcal{N}, \mathcal{R}, \mathcal{D}\}$ is the set of *observable outcomes* in the stage game, comprised of the events: no loan, repayment and default. Involuntary and voluntary defaults cannot be distinguished by anyone other than the borrower concerned, and are both denoted by \mathcal{D} .

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. We focus on equilibria where all lenders play the same strategy, and all borrowers play the same strategy. We would also like our equilibria to be robust. A strong notion of robustness is sequential strictness:

Definition 1 *An equilibrium σ of Γ^∞ is sequentially strict if every player has strict incentives to play her equilibrium action at every information set, whether this information set arises on or off the equilibrium path.*

¹⁰Incentives for the lender have to be provided within the period. This is the case even if borrowers are provided information on the past behaviour of lenders — see Bhaskar and Thomas (2018).

Sequential strictness is a demanding requirement, possibly too demanding, as it rules out any equilibrium in mixed strategies. A weaker criterion is to require sequential strictness in a “nearby game”. If one takes the view that agents always face random payoff shocks, and that the unperturbed game is an idealisation, then sequential strictness in the nearby game is an adequate robustness criterion. When this is met, the equilibrium is said to be *purifiable*, as in Harsanyi (1973). In our game, purifiability excludes belief-free equilibria, where a player is indifferent between two mixed actions α and α' , at two histories h and h' , but plays α at h and α' at h' . With payoff shocks, her indifference is broken in the perturbed game, and her behaviour must be the same at both histories.¹¹

The remainder of this section makes the robustness criterion precise. Define $\Gamma(\varepsilon)$, a perturbed version of the stage game Γ , indexed by $\varepsilon > 0$, a scaling parameter. Let X denote the set of player decision nodes in Γ . At each node $x \in X$, the payoff of the player who moves from one of her (two) actions is augmented by εz_x , where z_x is the realisation of a random variable Z_x with bounded support. The random variables $\{Z(x)\}_{x \in X}$ are independently distributed, and their distributions are atomless. The player who moves at node x observes the realisation z_x before she moves. In the repeated version of the perturbed game, $\Gamma^\infty(\varepsilon)$, we assume that the shocks for any player are independently distributed across periods.¹² In the lender-borrower interpretation of the trust game, the lender gets an idiosyncratic payoff shock from making a loan, while the borrower gets an idiosyncratic shock from wilfully defaulting. An equilibrium σ of Γ^∞ is *purifiable* if there exists a sequence of equilibria $\sigma(\varepsilon_n)$ of $\Gamma^\infty(\varepsilon)$ that converges to σ for any strictly positive sequence $\varepsilon_n \rightarrow 0$.

The following lemma, proved in Appendix A.5.3, shows that purifiability is a generalisation of sequential strictness.

Lemma 1 *Every sequentially strict equilibrium of Γ^∞ is purifiable.*

3 Benchmarks

3.1 Infinite Memory

Suppose that each lender can observe the entire history of outcomes in \mathcal{O} of each borrower he is matched with, and that the borrower observes no information about the lender. Suppose

¹¹Bhaskar, Mailath, and Morris (2013) show that such conditioning on history is not purifiable in a large class of stochastic games.

¹²The assumption that the lender’s shocks are independently distributed across periods is not essential.

payoff parameters are such that there exists an equilibrium where lending takes place.¹³

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 $(1 - \delta)g$.¹⁴ 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

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

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

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

$$V^K(K) = \delta^K V^K(0).$$

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 (2) with $K = \bar{K}$. We evaluate the payoffs of a 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. $\bar{W} = \frac{1}{1 + \lambda\bar{K}}$.

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.1 we show that the highest payoff the borrower can achieve in

¹³That 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 $g < \frac{\delta(1-\lambda)}{1-\delta(1-\lambda)}$.

¹⁴Per-period payoffs are normalised by multiplying by $(1 - \delta)$.

any equilibrium is given by:¹⁵

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

To sustain the equilibrium payoff V^* , we assume that players observe the realisation of a public randomisation device at the beginning of each period, and that past realisations of the randomisation device are also a part of the public history. The payoff V^* can be achieved by the borrower being excluded for \bar{K} periods with probability x^* , and for $\bar{K} - 1$ periods with probability $1 - x^*$. This gives rise to a steady-state proportion of borrowers in good standing equal to $\frac{1}{1 + \lambda(\bar{K} - 1 + x^*)}$. This equals the lender's expected payoff, W^* .

To summarise: \bar{V} and \bar{W} are the *constrained efficient payoffs* for the borrower and lender respectively, that reflect both the integer constraint and the borrower's incentive constraint under imperfect monitoring. V^* and W^* are the *fully efficient payoffs* — these reflect 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.

3.2 Perfect Bounded Memory

Henceforth, we shall assume bounded memory: 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. (The borrower does not observe any information regarding the lender.)

Our first result is a negative one — if the lender has full information regarding the past K outcomes of the borrower, then no lending can be supported.

Proposition 1 *Suppose that K is arbitrary and the lender observes the finest possible partition of \mathcal{O}^K , or $K = 1$ and the information partition is arbitrary. If the equilibrium is sequentially strict or purifiable, the lender never lends and the borrower never repays.*

The proof is an adaptation of the argument in Bhaskar, Mailath, and Morris (2013). To get some intuition, suppose that $K \geq \bar{K}$ and the information partition is the finest possible, and consider a candidate equilibrium where a borrower is lent to unless her record has any instance of \mathcal{D} in the last K periods. A borrower with a clean record prefers to repay. Consider a borrower with exactly one default which occurred exactly K periods ago. She has

¹⁵In 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 randomisation device that allows her to wilfully default with some probability, and such defaults are not punished. Furthermore, past realisations of the randomisation device must also be a part of the infinite public history. The assumption that mixed strategies are observable seems strong and unrealistic.

incentives identical to those of a borrower with a clean record, and will 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.

More generally, in a sequentially strict equilibrium, there cannot be any conditioning on a borrower's history. To show this, we first argue that a borrower who is given a loan cannot condition her repayment decision on ω^K , the outcome exactly K periods ago. Indeed, the lender tomorrow (and every subsequent lender) cannot observe ω^K , since it will disappear from the record. Consequently, the borrower's continuation value tomorrow cannot depend upon ω^K . Now compare two K -period histories of the borrower, h and h' , that are identical except with regard to ω^K . If the borrower plays different actions at these histories, then she must get the same flow payoff from both these actions at h and at h' (since, we just argued, she cannot be compensated via continuation payoffs). But this contradicts our assumption that the equilibrium was sequentially strict. Thus, the borrower cannot condition her behaviour on ω^K , and must take the same action at h and h' . Now if this action is to repay, then the lender must lend at both h and h' , while if it is to default, the lender must not lend at both histories. In either case, the lender cannot condition his behaviour on ω^K either, and must play the same action at both h and h' .

Having established that neither lender nor borrower can condition their behaviour today on ω^K , we can repeat the same argument to show that a borrower will not condition her behaviour on ω^{K-1} , the outcome $K - 1$ periods ago. An induction argument implies that neither borrower nor lender will condition their behaviour on any of the borrower's history, implying that the only sustainable outcome corresponds to the backwards-induction profile being played at every period.

The weaker requirement of purification ensures that even if an equilibrium of the unperturbed game is not sequentially strict, it must be sequentially strict in the perturbed game. (A player can be indifferent between two actions only on a measure zero set, and her behaviour on a negligible set is of no consequence to other players.) Thus the argument made in the previous paragraphs applies in *any* perturbed game. As a result, although there are belief-free equilibria in the unperturbed game that sustain lending, these have no counterpart when there are payoff shocks, and are therefore not purifiable.

The second part of the proposition, that there can be no conditioning on history if $K = 1$, applies to any information structure. To support lending, we will need $K \geq 2$. As we shall see, total breakdown of lending does not occur under all information structures. Surprisingly, *less information* may support cooperative outcomes. Although it remains the case that a

borrower whose most recent default is exactly K periods ago will repay a loan, a coarser information structure prevents the lender from targeting her, so the induction argument underlying the above proposition does not apply.

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.¹⁶ 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. At a general level, ours is an instance of information design (e.g. Kamenica and Gentzkow (2011)) although our methods are very different from the approach taken in this literature.

Let K denote the bound on memory chosen by the designer — we allow K to be arbitrarily large but finite. An information system provides information to the lender based on the past K outcomes in \mathcal{O} of the borrower. We assume that the borrower does not receive information on the past outcomes of the lender¹⁷ Information structures fall into two broad categories. A *deterministic* information (or signal) structure consists of a finite signal space S and a mapping $\tau : \mathcal{O}^K \rightarrow S$. 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 $\tau : \mathcal{O}^K \rightarrow \Delta(S)$.

Note that in both cases, the signal *does not* depend on past signal realisations, since otherwise one could smuggle in infinite memory on outcomes. We focus on partitional information structures since the efficiency gains from random information are restricted to overcoming integer problems.

¹⁶Thus, the designer cannot dictate the actions to be taken by any agent, and in particular cannot direct lenders to refrain from lending to defaulters.

¹⁷In Section 7.4 of Bhaskar and Thomas (2018) we show that such information would be useless, since no borrower would condition on it.

4.1 A Simple Information Structure

Since the length of memory, K , can be chosen without constraints, we assume that $K \geq \max\{\bar{K}, 2\}$.¹⁸ The information structure is given by the following binary partition of \mathcal{O}^K , which we call the *simple* information structure. 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. This section shows that lending can be supported under the simple information structure.

The borrower has complete knowledge of her own private history. 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 suffices to describe the borrower’s incentives. Partition the set of private histories into $K + 1$ equivalence classes, indexed by $m := \min\{K + 1 + t' - t, 0\}$, where t denotes the current period, and t' denotes the date of the most recent incidence of \mathcal{D} in the borrower’s history. 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 signal. When $m \geq 1$, this value is the borrower’s private information. In particular, among borrowers 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

$$V^K(m) = \delta^m V^K(0), \quad m \in \{1, \dots, K\}.$$

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 are summarised by m . The borrower’s incentives at $m = 1$ are identical to those at $m = 0$ — for both types

¹⁸If actual memory \bar{K} is greater than K , it is straightforward to reduce its effective length to 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 \bar{K} periods ago — this follows from arguments similar to those underlying Proposition 1.

of borrower, their current action has identical effects on their future signal. Therefore, a borrower of type $m = 1$ will always repay. 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, i.e. we require

$$(1 - \delta)g > \delta(1 - \lambda) [V^K(K - 1) - V^K(K)]. \quad (3)$$

The inequality above is satisfied whenever $K \geq \bar{K}$ (see Appendix A.3). When $m = K$, a borrower who repays still faces $K - 1$ periods of exclusion, and so a new default is only punished by one additional period of exclusion. But we assume ($\bar{K} \geq 2$) that at least two periods of punishment are required to incentivise repayment.

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

$$(1 - \delta)g - \delta(1 - \lambda)[V^K(m - 1) - V^K(K)] = (1 - \delta)g - (1 - \lambda)(\delta^m - \delta^{K+1})V^K(0). \quad (4)$$

The right-hand side of the above expression is defined for all real valued m ; we have seen that when $m = K$, it is positive, while when $m = 1$, it is negative. Thus there exists a real number, denoted $m^\dagger(K) \in (1, K)$, that sets the payoff difference equal to zero. For generic payoffs $m^\dagger(K)$ is not an integer, and we assume this to be the case, ensuring that the equilibrium is sequentially strict. Let $m^*(K) := \lfloor m^\dagger(K) \rfloor$ denote 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 her sentence will be on her best behaviour.

Since the lender has imperfect information regarding the borrower's K -period history, we compute the lender's beliefs about those histories using Bayes' rule. We 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.7 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 differ from the steady state one. In every period, the probability of involuntary default is constant, and equals λ . Under our strategy profile, a borrower with a bad credit history never gets a loan and hence transits deterministically through the states $m = K, K - 1, \dots, 1$. Therefore, the invariant distribution over values of m induced by this strategy profile gives equal probability to each of these states. Consequently,

the lender attributes probability $\frac{m^*(K)}{K}$ to a borrower with signal B repaying a loan. Simple algebra shows that lending to a borrower with a bad credit history is strictly unprofitable for the lender if

$$\frac{m^*(K)}{K} < \frac{\ell}{1 + \ell}. \quad (5)$$

Suppose that ℓ is large enough that the lender’s incentive constraint (5) is satisfied. Then he finds it strictly optimal not to lend after B , and to lend after G . A borrower who is granted a loan has strict incentives to repay as long as $m \leq m^*(K)$ and to default if $m > m^*(K)$. Thus, there exists an equilibrium that is sequentially strict (and 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.

It remains to identify conditions on the parameters ensuring that the incentive constraint (5) is satisfied. In Appendix A.3, we show that $m^*(K)$ is bounded. When K becomes large enough, further increments in K have negligible effects on $V^K(0)$. Consequently, m^* (the maximal length of *remaining* punishment such that re-offending is unprofitable) becomes independent of K . Consequently, $\frac{m^*(K)}{K} \rightarrow 0$ as $K \rightarrow \infty$, giving the following proposition:

Proposition 2 *A sequentially strict equilibrium, where the lender lends after observing G and does not lend after observing B , exists as long as K is sufficiently large.*

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. By giving the lender coarse information, his incentive constraints are pooled, and need to only hold on average. With fine information, the lender’s incentive constraint may be violated only for one type, but this results in unravelling and a total breakdown, as in Proposition 1. In other words, coarse information endogenously generates borrower adverse selection that mitigates *lender moral hazard*.¹⁹

¹⁹It is well known that exogenous adverse selection can help solve moral hazard problems — see e.g. Ghosh and Ray (1996).

5 Efficient Equilibria

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 the borrower's constrained efficient payoff \bar{V} can be achieved for all parameter values. When lending is very profitable (ℓ is small), the equilibrium has lenders making loans with positive probability to bad credit risks, resulting in low profits for lenders.

5.1 Pure strategy equilibrium when ℓ is large

Consider the pure strategy profile set out in the previous section, with \bar{K} memory. When punishments are of the minimal length \bar{K} , any punishment whose effective length is shorter will not incentivise repayment. Consequently, $m^*(\bar{K}) = 1$, 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

$$\frac{1}{\bar{K}} < \frac{\ell}{1 + \ell}$$

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 achieve the fully efficient payoffs V^* and W^* by using a random signal structure. 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 B is observed with probability x and G with probability $1 - x$. Let $x > x^*$, 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 $m^* = 1$ under this random signal structure, so that the repayment probability at signal B remains low enough that 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 the random version of the simple information structure.*

Adding moderate exogenous adverse selection among borrowers in our model may actually improve matters, by relaxing the lenders' constraint. Suppose that there are two types of borrowers, who differ only in their rates of involuntary default. Riskier borrowers will be over-represented amongst the borrowers with a B signal, reducing the lender's incentive to lend at B . Thus, exogenous borrower adverse selection alleviates the lenders' tendency for rogue lending, but it is not essential.

Aggregate shocks may result in a divergence from the steady state, affecting lender incentives. Consider a large, unanticipated, temporary increase in λ (the rate of involuntary default) in period t , that raises this cohort's subsequent share among bad credit risks. In period $t + \bar{K}$, lenders are aware that a larger than usual share of bad credit risks have an incentive to repay their loans, and it may become profitable to lend. If this is anticipated by borrowers, then voluntary default becomes profitable in period $t + 1$, breaking the equilibrium. An information policy of "forgiving" a proportion of the excess defaults, giving those defaulters a clean record, solves the problem. This will not affect the incentives of period t borrowers, provided they do not observe the aggregate shock contemporaneously.

5.2 Mixed strategy equilibrium when ℓ is small

Consider now the case where $\ell < \frac{1}{\bar{K}-1}$, and suppose that lenders lend with positive probability upon observing B . This permits an equilibrium where the length of exclusion after a default is no greater than \bar{K} — the effective length is strictly less, since exclusion is probabilistic. This may appear surprising — if a lender is required to randomise 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 behaviour of the population of lenders changes the mix of different types of borrower among those with signal B , raising the proportion of those with larger values of m . This raises the default probability of borrowers with a bad credit history, and disciplines lenders.

Consider the strategy profile with \bar{K} memory where the lender always extends a loan at signal G , and with probability p at signal B , and a borrower with $m > 1$ never repays the loan, and repays with probability q if $m = 0$ or $m = 1$. 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, 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 she is indifferent between repaying and defaulting. 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

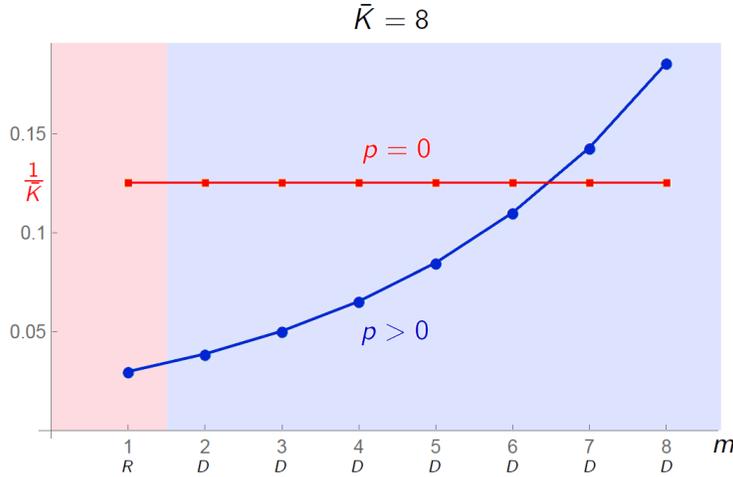


Figure 2: *Stationary probabilities conditional on signal B: $\mu_m/(1-\mu_0)$ for $m = 1, \dots, \bar{K}$. Illustrated for $p = 0$ and $p = 0.23$. (For $q = 1$.)*

appeal of defaulting, so a borrower with $m > 1$ will continue to default when $p > 0$.

Now consider the incentives of the lender at signal B , and let $\pi(p, q)$ denote the probability with which he expects a loan to be repaid at B . The lender is indifferent between lending and not lending at B if and only if $\pi(p, q) = \frac{\ell}{1+\ell}$. Loans are repaid only by type $m = 1$ (and only with probability q). Under the mixed strategy profile, the fraction of those types among the pool of B -signal borrowers is less than $1/\bar{K}$. This is because, under the mixed profile, a borrower with $m > 1$ gets a loan with positive probability, defaults, and restarts her punishment phase. Thus, the steady-state distribution $\{\mu_m(p, q)\}_{m=0}^{\bar{K}}$ puts less weight on lower values of m , as illustrated in Figure 2.

When $q = 1$, the probability that a loan made at signal B is repaid is

$$\pi(p, 1) := \frac{\mu_1(p, 1)}{1 - \mu_0(p, 1)}. \quad (6)$$

In Appendix A.5.1 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 $\pi(\tilde{p}, 1) \leq \frac{\ell}{1+\ell}$, then 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 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.²⁰ We focus on this equilibrium. The probability that a loan made at history B is repaid is now

$$\pi(\tilde{p}, q) := \frac{q \mu_1(\tilde{p}, q)}{1 - \mu_0(\tilde{p}, q)} = q \pi(\tilde{p}, 1). \quad (7)$$

We establish the second equality in Appendix A.5.2. Clearly, $\pi(\tilde{p}, q)$ is a continuous, strictly increasing function of q . Since we are considering the case where $\pi(\tilde{p}, 1) > \frac{\ell}{1+\ell}$, and since $\pi(\tilde{p}, 0) = 0$, there exists a value \tilde{q} setting the repayment probability $\pi(\tilde{p}, \tilde{q})$ equal to $\frac{\ell}{1+\ell}$.

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 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 idea underlying our mixed equilibrium is reminiscent of the work of Kandori (1992) and Ellison (1994) on the prisoner's dilemma played in a random matching environment. Ellison shows that when ℓ is small, punishments must be finely tuned, possibly using a public randomisation device. They must be severe enough so that a player does not want to start the contagion process, but not so severe that she is unwilling to join in once it begins. Here, when ℓ is small, mixing plays a similar role, by raising the proportion of defaulters among bad credit risks.

The borrower payoff in the mixed equilibrium lies between \bar{V} and V^* . It is strictly greater than \bar{V} since her effective punishment phase is at most \bar{K} periods. When ℓ is so low that the borrower also mixes, then she gets the payoff V^* — her incentive constraint is satisfied with equality at G . 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 he only makes positive profits when lending to a borrower with signal G , 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 signal G with some probability, the lender's profits fall further, and as ℓ tends to zero, so do profits.

²⁰Appendix A.5.3 proves that the mixed equilibrium where $m \in \{0, 1\}$ repay with the same probability is purifiable. It also shows that there are other, non-purifiable equilibria, that may be more efficient.

An Example: We consider parameter values such that $\bar{K} = 4$.²¹ Consider the pure strategy profile when $K = \bar{K}$, where the lender extends a loan only after G . The invariant distribution over m -values is uniform, and the lender is repaid after lending at B with probability $\frac{1}{4}$. 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 4-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 $m^*(K) = 7$. Under the invariant distribution over borrower types, only a quarter of the population with a clean history, and 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 a clean history is 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 $m \in \{0, 1\}$.²² 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 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. \square

²¹Specifically, $\delta = 0.9$, $\lambda = 0.1$ and $g = 2$.

²²If $\ell = 0.1$, $K = 89$ periods of exclusion are needed to support a pure strategy equilibrium. In this case $m^*(K) = 8$. The lender's payoff is 0.1 and the borrower's payoff at a clean history is 0.526.

5.3 Initiating Trust

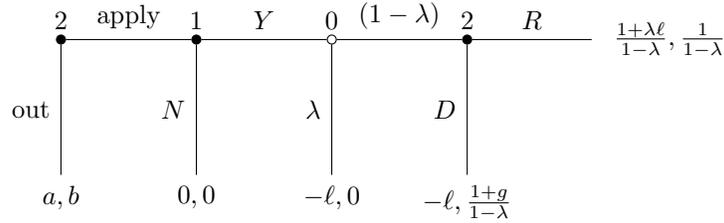


Figure 3: *Trust game with costly initiation.*

Suppose the interaction needs to be initiated by the borrower, i.e. the player who must be trusted. This modification of the trust game is illustrated in Figure 3. First, the borrower (player 2) must decide whether or not to apply for a loan, where application entails a small cost. 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*. Assume the simple information structure with \bar{K} periods memory.²³

Under this information structure, prior application transforms the interaction between an individual lender and borrower into a signalling 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 signal G apply, are given a loan and repay. Borrowers with signal 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 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 5 *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 supported by beliefs assigning probability one to an applicant with a bad credit history defaulting. The borrower beliefs satisfy the D1 criterion. The equilibrium payoffs approximate V^* and W^**

²³For the full efficiency result, we use the random version of the simple information structure, under which the borrower gets a good signal if she has not defaulted in the last \bar{K} periods, and if she has defaulted exactly once \bar{K} periods ago, in which case she gets signal G with probability $1 - x$.

We conclude that when the interaction must be initiated, at a small cost, by the party to be trusted, fully efficient payoffs can be sustained for both parties.

6 Extensions

6.1 Screening Borrowers

Fix the simple information structure with memory length K , and suppose that parameter values are such that there exists a pure strategy equilibrium, as in Section 4. We now ask: can a single lender, who meets a borrower with a bad credit history, deviate by offering an alternative contract? In particular, can he offer a contract that is acceptable only to good credit risks, and unacceptable to bad credit risks? In Appendix A.8, we show that the answer is no. Any contract that is acceptable to a borrower of type m will be strictly better than refusal for a borrower with type $m' > m$. The intuition for this is straightforward. If a borrower of type m intends to default, then her overall value from acceptance, relative to refusal, is increasing in m . Similarly, if the borrower intends to repay, then her overall value, relative to refusal, is also increasing in m , due to the possibility of involuntary default. Thus the value of the borrower from *any* contract offered by the deviating lender is increasing in m , implying that the separation desired by the lender cannot be achieved.

Nevertheless, the deviating lender might offer an alternative contract that uniformly reduces default rates, without inducing separation. For example, the lender might offer a smaller loan at a B signal, that requires a smaller gross repayment, thereby reducing the borrower's incentive to default. Smaller loan sizes induce more borrower types to repay. The deviating lender's profits from making this loan equals the gross repayment times the repayment probability, minus the sum lent and the fixed cost of lending. If the fixed costs of lending are sufficiently large, then reducing loan size may not be profitable, even though this raises the repayment probability.²⁴

Finally, let us suppose that the deviating lender offers a contract with deferred repayment — repayment is due in the next period, rather than at the end of the current period, and has the same present value. This has subtle effects, since borrowers are closer to the end of the punishment phase when they are required to repay. In Appendix A.8, we show that deferred repayment may affect the default behaviour of two types of borrower — types $m^* + 1$ and $m = 1$ — but leaves the decisions of other borrower types unchanged. Type $m^* + 1$ may

²⁴In real-world credit markets, borrowers in bad standing are often not entirely excluded, but are granted limited access to credit. Thus the punishment for bankruptcy is not total exclusion but credit limitation.

or may not switch her behaviour from defaulting to repaying — the effect is ambiguous. On the other hand, type $m = 1$ will switch to defaulting, since in the next period she get another loan and has to make two repayments, doubling her incentive to default. Since types are uniformly distributed, conditional on a B signal, offering deferred payments either reduces the overall repayment probability, or leaves it unaffected, so that the deviation is not profitable. These extensions suggest that our equilibrium is robust to allowing for deviations to alternative contracts.

6.2 Empirical Evidence and Implications

Bounded memory is legally enforced in credit markets, as bankruptcy flags have to be removed from borrowers' records after a fixed length of time. Empirical evidence²⁵ shows that consumers experience a jump in credit scores in the quarter of removal of the bankruptcy flag, which leads to a large increase in their credit limits and borrowing. For example, Dobbie et al. find that the increase in credit scores corresponds to an implied 3 percentage point reduced default risk, on a pre-flag removal risk of 32 percent. It is implausible that such a large change can be attributed entirely to a sudden shift in consumer behaviour. Rather, it reflects a change in lenders' perceptions, based on available information. This shows that memory constraints are real, and information leaves the market when flags are dropped.

Our theory predicts unravelling: lenders should actively target those customers whose flag is about to disappear. Existing empirical work is silent on unravelling, since it has focused entirely on the comparison before vs. after flag removal, and does not examine the dynamic path of lending prior to removal. This is a fruitful area for future empirical research.

The degree of unravelling should increase when default flags provide little additional information about a consumer's underlying propensity to default. There is no consensus on the matter. Musto and Dobbie et al. find that removed flags are predictive of future credit delinquency, even after conditioning on other available information, in particular a consumer's credit score, while Gross et al. find no predictive effect. With improvements in information-processing and the rise of big data, one would expect that the incremental informativeness of default flags would be reduced, increasing the tendency towards unravelling.

To summarise, our paper suggests directions for new empirical work on consumer credit. What is the trajectory of credit limits and loan terms for borrowers prior to bankruptcy flag removal, and how do these compare with terms for borrowers whose flags have just been

²⁵See Musto (2004), Gross, Notowidigdo, and Wang (2016) and Dobbie, Goldsmith-Pinkham, Mahoney, and Song (2016).

removed? How likely are such borrowers to be delinquent, compared to similar borrowers without past defaults, and compared to borrowers whose flags have just been removed? Finally, what are the social welfare consequences of existing credit scoring systems? The stated purpose of credit rating agencies is to help individual lenders predict delinquency — not to sustain socially efficient outcomes. These two goals may well be in conflict.

6.3 Generalising our Results

Our results can be generalised to sustain a Pareto-efficient outcome when the stage game is a two-player game of perfect information where moral hazard is effectively one-sided. (See Bhaskar and Thomas (2018) for details.) Consider a Pareto-efficient terminal node, z^* , that strictly Pareto-dominates the backwards induction outcome. Let both players conjecture that any deviation from the path to z^* is followed by players continuing with the backwards induction strategies. Suppose that only one player has an incentive to deviate from the path to z^* given this conjecture about continuation play, and that this incentive arises at a single node. The analysis of this paper can be extended to show that we can sustain the outcome z^* , by providing coarse information about the player who has an incentive to deviate.

The class of stage-games that our results extend to includes, in particular, all games where each player moves at most once along any path of play, as only the second mover can have an incentive to deviate. To illustrate, consider the prisoner’s dilemma with sequential moves, where the efficient outcome is mutual cooperation. Player 1 cannot profitably defect, since player 2 can respond by defecting. So only player 2 move has an incentive to deviate in the one-shot game. But we can also accommodate some games where the players move more than once. In the centipede game in Figure 4, if the players play “across” at each node, this results in payoffs (3, 3). If both players expect this outcome, then only player 2 has an incentive to deviate. Since moral hazard is one-sided, the analysis of this paper applies. However, if player 1’s payoff from playing d_3 is changed from 2 to 4, then he would also have an incentive to deviate at this node, and our analysis would not apply.

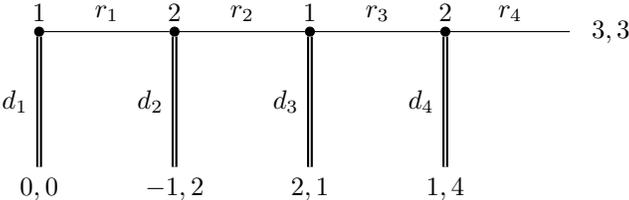


Figure 4: Centipede game.

7 Conclusion

Existing work on community enforcement has focused entirely on simultaneous-move games. Many economic interactions have a sequential structure, the trust game being a paradigmatic example. Although moral hazard is one-sided, incentivising the side that is not subject to moral hazard to discipline 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.

Methodologically, our analysis differs from much of the recent work on repeated games that employs belief-free equilibria. We have insisted on equilibria where players have strict incentives, at least when they are subject to payoff shocks. The information structures and equilibria that sustain efficient outcomes are robust, simple and intuitive.

A Appendix: For Online Publication

A.1 Proofs Related to Section 3.1

We derive the upper bound on the borrower's equilibrium value under infinite memory when the borrower's mixed strategies are not observable. 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:

$$V^* \leq (1 - \delta) + \delta [\lambda V^P + (1 - \lambda) V^*]. \quad (\text{A.1})$$

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 be optimal in the current period, we have the following necessary incentive constraint

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

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

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

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

$$V^* \leq 1 - \frac{\lambda}{1 - \lambda}g. \quad (\text{A.3})$$

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, which implies that inequality A.3 holds with equality.

A.2 Proof of Proposition 1

Fix a sequentially strict equilibrium of the game with perfect K period memory, and let σ denote the pure strategy of an arbitrary borrower and let ρ denote the pure strategy of an

arbitrary lender. We do not assume that all lenders/borrowers follow the same strategy, but to avoid needless notation, we do not index these strategies by individual identities. Since lenders have K period memory, $\rho : \mathcal{O}^K \rightarrow \{Y, N\}$.²⁶ Since the borrower observes her entire history, her strategy is a sequence $(\sigma_t)_{t=1}^\infty$, where $\sigma_t : \mathcal{O}^{t-1} \rightarrow \{R, D\}$. For an arbitrary natural number T and any integer j , $0 \leq j \leq T$, define the equivalence relation \sim_j on \mathcal{O}^T such that $h \sim_j h'$ iff the last j outcomes in h and h' are identical.

The proof is by induction, and consists of the following steps:

1. Any lender's strategy is measurable with respect to \sim_K , due to the bound on memory.
2. If each lender's strategy is measurable with respect to \sim_j for some integer j , then any borrower's strategy is measurable with respect to \sim_{j-1} .
3. If each borrower's strategy is measurable with respect to \sim_j , with $j \leq K$, then any lender's strategy is also measurable with respect to \sim_j .

These steps imply that every lender strategy and every borrower strategy is measurable with respect to \sim_0 , so that there cannot be any conditioning on borrower histories. Thus, the backwards induction outcome of the trust game must be played at every history.

Point 1 is immediate, so we turn to 2. Assume that strategy of every lender is measurable with respect to \sim_j . Consider a borrower at date t with a history h , who has been made a loan. Her continuation value at date $t+1$ depends on $(h, \omega_t) \in \mathcal{O}^t$, where $\omega_t \in \{\mathcal{D}, \mathcal{R}\}$ is the outcome at date t . Since the strategy of any lender at date $t+1$ or any subsequent date is assumed to be measurable with respect to \sim_j , the borrower's continuation value at date $t+1$, $V_{t+1}(h, \omega_t)$, is measurable with respect to \sim_j . Thus, if $h \sim_{j-1} h'$, $V_{t+1}(h, \omega_t) = V_{t+1}(h', \omega_t)$. We now show that if the borrower has strict best responses at h and h' , with $h \sim_{j-1} h'$, then $\sigma_t(h) = \sigma_t(h')$. Suppose that the borrower plays differently at the two histories, e.g. repays at h and defaults at h' . Since repaying is optimal at h ,

$$(1 - \delta)g \leq \delta(1 - \lambda)[V_{t+1}(h, \mathcal{R}) - V_{t+1}(h, \mathcal{D})].$$

Since defaulting is optimal at h' ,

$$(1 - \delta)g \geq \delta(1 - \lambda)[V_{t+1}(h', \mathcal{R}) - V_{t+1}(h', \mathcal{D})].$$

However, since $V_{t+1}(h, \omega_t)$ is measurable with respect to \sim_j , $V_{t+1}(h, \mathcal{D}) = V_{t+1}(h', \mathcal{D})$ and

²⁶We can augment \mathcal{O} with a "dummy" outcome to describe histories in the initial K periods of the game.

$V_{t+1}(h, \mathcal{R}) = V_{t+1}(h', \mathcal{R})$, which implies:

$$(1 - \delta)g = \delta(1 - \lambda)[V_{t+1}(h, \mathcal{R}) - V_{t+1}(h, \mathcal{D})].$$

Thus the equilibrium cannot be sequentially strict if the borrower plays different actions at two $(j - 1)$ -equivalent histories. A contradiction. This proves point 2.

Turning to point 3, if the current borrower is playing a pure strategy that is measurable with respect to \sim_j , and $j \leq K$, then the lender has point beliefs about the behaviour of the borrower — that she will either default for sure or repay for sure. Thus the lender's strategy is also measurable with respect to \sim_j . This completes the proof for any sequentially strict equilibrium of Γ^∞ .

Now let us consider the perturbed game. The payoff to the lender from lending is augmented by εy , where y is the realisation 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 realisation of a random variable that is distributed on $[0, 1]$ with a continuous cumulative distribution function, F_Z .

A strategy for the lender is now $\rho : \mathcal{O}^K \times [0, 1] \rightarrow \{Y, N\}$, while that of the borrower is a sequence $(\sigma_t)_{t=1}^\infty$, where $\sigma_t : \mathcal{O}^t \times [0, 1] \rightarrow \{R, D\}$. We modify the definition of measurability as follows:

- σ is measurable with respect to \sim_j if for any t , whenever $h \sim_j h'$, $\sigma_t(h, z) = \sigma_t(h', z)$, except possibly on a set of z -values that has F_Z -measure zero.
- ρ is measurable with respect to \sim_j if for any t , whenever $h \sim_j h'$, $\rho(h, y) = \rho(h', y)$, except possibly on set of y -values that has F_Y -measure zero.

The proof is by induction, and the three steps are exactly as before. Since the proof of the first step is immediate, we turn to the second. Assume that the strategy of every lender is measurable with respect to \sim_j . Fix a value of z such that the borrower plays differently at h and h' , where $h \sim_{j-1} h'$; she repays at h and defaults at h' . By the same arguments as in the previous proof, we deduce that

$$(1 - \delta)(g + \varepsilon z) = \delta(1 - \lambda)[\tilde{V}_{t+1}(h, \mathcal{R}) - \tilde{V}_{t+1}(h, \mathcal{D})], \quad (\text{A.4})$$

where $\tilde{V}_{t+1}(h, \omega)$ denotes the borrower's ex-ante value function at history (h, ω) , before her payoff shock in period $t + 1$ is realised. However, since the left-hand side of (A.4) is strictly

increasing in z , (A.4) can hold for at most one value of z in $[0, 1]$. This establishes step 2 in the argument.

Turning to step 3, assume that the strategy of the borrower is measurable with respect to \sim_j , with $j \leq K$, and let $h \sim_j h'$. Let $z^*(h)$ denote the value of z such that the borrower is indifferent between repaying and defaulting at h , with $z^*(h) = 0$ (resp. 1) if the borrower always prefers to default (resp. repay). Thus the payoff to the lender from lending at h is

$$F_Z(z^*(h))(1 + \varepsilon y) - [1 - F_Z(z^*(h))]\ell.$$

Since the payoff from lending is increasing in y , the lender's optimal strategy is characterised by a threshold, $y^*(h)$, the only point at which he is possibly indifferent between lending and not lending. Further, since the borrower strategy is measurable with respect to \sim_j , $z^*(h') = z^*(h)$, and the lender's optimal threshold at h' must equal $y^*(h)$. Thus the lender's strategy is measurable with respect to \sim_j . This completes the proof of Proposition 1.

A.3 Proofs Related to Sections 4.1

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 (3)):

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

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 the following, stronger result: 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

$$(1 - \delta)g > \delta(1 - \lambda) \left(V^{\bar{K}}(1) - V^{\bar{K}}(\bar{K}) \right) = \delta(1 - \lambda)(\delta - \delta^{\bar{K}})V^{\bar{K}}(0). \quad (\text{A.5})$$

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.5), 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 (3).

The key step in proving Proposition 2 is showing that $m^\dagger(K)$ is bounded. m^\dagger is the real value of m that sets the right-hand side of (4) equal to zero, and is given by:

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

Since

$$\begin{aligned} \lim_{K \rightarrow \infty} V^K(0) &= \frac{(1-\delta)}{1-\delta(1-\lambda)}, \\ \lim_{K \rightarrow \infty} m^\dagger(K) &= \ln \left[\frac{[1-\delta(1-\lambda)]g}{(1-\lambda)} \right] \frac{1}{\ln \delta}, \end{aligned}$$

which proves the required result.

A.4 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} periods and with complementary probability for $\bar{K}-1$ 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

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

The borrower agrees to repay at G if and only if

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

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

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

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 $x \in [x^*(\bar{K}), 1]$.²⁷

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

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

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

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

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

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

and that

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

²⁷In the non-generic case where \bar{K} satisfies (1) with equality, we have that $x^*(\bar{K}) = 1$.

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 Proofs Related to Section 5.2

We now prove claims related to the mixed strategy equilibrium.

A.5.1 First Claim

We show that $\pi(p, 1)$, defined in (6), is a strictly decreasing function of p . (Continuity for $p \in [0, 1]$ is immediate.)

We first derive the invariant distribution. For any given $K \geq 2$, 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 $\{0, 1, 2, \dots, \bar{K}\}$. 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

$$\mu_{m-1} = (1 - p)\mu_m \quad \text{if } m > 1. \quad (\text{A.11})$$

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

$$\mu_{\bar{K}} = \lambda(\mu_0 + p\mu_1) + p \sum_{m=2}^{\bar{K}} \mu_m. \quad (\text{A.12})$$

Finally, a borrower with $m = 1$ transits to $m = 0$ unless she gets a loan and suffers involuntary

default, and so

$$\mu_0 = (1 - \lambda)\mu_0 + (1 - p\lambda)\mu_1. \quad (\text{A.13})$$

Finally, we must have $\sum_{m=0}^K \mu_m = 1$. Solving the above system, we obtain

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

so that

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

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

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

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.5.2 Second Claim

We now establish the second equality in (7). For any given $K \geq 2$ and $(p, q) \in (0, 1)^2$, the invariant distribution $\mu(p, q)$, is given by

$$\begin{aligned} \mu_0 &= q(1 - \lambda)\mu_0 + (1 - p(1 - q(1 - \lambda)))\mu_1, \\ \mu_m &= (1 - p)^{K-m}\mu_K, \quad 1 \leq m \leq K, \\ \mu_K &= (1 - q(1 - \lambda))\mu_0 + p(1 - q(1 - \lambda))\mu_1 + p \sum_{m=2}^K \mu_m, \end{aligned}$$

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 (7).

A.5.3 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 realisation 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 realisation of a random variable that is distributed on $[0, 1]$ with a continuous cumulative distribution function, F_Z .

The proof of Lemma 1 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 realisations 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 behaviour 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. It is worth noting that the equilibrium singled out in the proposition is not *regular* — it is not isolated, since it is one of a continuum of equilibria. First, we explain 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 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.

Turning to the proof, observe that 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.²⁸

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 realisation 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

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

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

$$\tilde{V}^{\bar{K}}(1, \bar{p}) = \bar{p}(1 - \delta) \left(1 + \varepsilon \int_{\bar{z}}^1 (z - \bar{z}) dF_Z(z) \right) + \delta \left[\bar{p} \lambda \tilde{V}^{\bar{K}}(\bar{K}, \bar{p}) + (1 - \bar{p} \lambda) \tilde{V}^{\bar{K}}(0, \bar{p}) \right]. \quad (\text{A.15})$$

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

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

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

$$(1 - \delta)(g + \varepsilon \bar{z}) - \delta(1 - \lambda) \left(\tilde{V}^{\bar{K}}(0, \bar{p}) - \tilde{V}^{\bar{K}}(\bar{K}, \bar{p}) \right) = 0. \quad (\text{A.17})$$

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

$$\frac{\bar{q} \mu_1(\bar{p}, \bar{q})}{1 - \mu_0(\bar{p}, \bar{q})} - \frac{\ell}{1 + \ell + \varepsilon \bar{y}} = 0. \quad (\text{A.18})$$

²⁸The standard proofs for purifiability, in Harsanyi (1973), Govindan, Reny, and Robson (2003) and Doraszelski and Escobar (2010), apply for regular equilibria.

When $\varepsilon = 0$, equations (A.17) and (A.18) 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.17) and (A.18) 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.17) and (A.18). 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 by²⁹

$$\gamma(p, q) = (1 - \delta)g - \delta(1 - \lambda)(V^{\bar{K}}(0, p, q) - V^{\bar{K}}(\bar{K}, p, q)) = 0. \quad (\text{A.19})$$

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

$$\phi(p, q) = q\mu_1(p, q)(1 + \ell) - \ell(1 - \mu_0(p, q)) = 0. \quad (\text{A.20})$$

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 $V^{\bar{K}}(\bar{K}, p, q)$ satisfy

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

Differentiating with respect to p , we obtain

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

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 $(p, q) = (\tilde{p}, \tilde{q})$.

²⁹Although, for every $m \in \{0, 1, \dots, \bar{K}\}$, the level of $V^{\bar{K}}(m, p)$ is independent of q when $p = \tilde{p}$, its slope is not. We therefore emphasise the dependence of $V^{\bar{K}}$ on q in the remainder of this section.

Differentiating ϕ with respect to q gives

$$\phi_q(p, q) = \frac{\partial \mu_1(p, q)}{\partial q} q(1 + \ell) + \mu_1(p, q)(1 + \ell) + \frac{\partial \mu_0(p, q)}{\partial q} \ell. \quad (\text{A.21})$$

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.21), we obtain

$$\phi_q(p, q) = \mu_1(p, q)(1 + \ell) \left[1 - \frac{QC}{(1 - Q)(A - QB)} \right] + \frac{(1 - \lambda)C}{(1 - Q)(A - QB)} (1 - \mu_0(p, q)) \ell. \quad (\text{A.22})$$

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.22) gives

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

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.6 Proof of Proposition 5

Proof. Under the random version of the simple information system, the borrower is excluded with probability x for $\bar{K} - 1$ periods and with complementary probability for \bar{K} periods. We can choose x strictly greater than $x^*(\bar{K})$ satisfying (A.7) (see the proof of Proposition 3), but sufficiently close to it. Then the borrower with a good credit record has strict incentives to repay, and her payoff is close to V^* . The average payoff of the lender is close to W^* , given that borrower exclusion is minimal.

It is straightforward to verify that the equilibrium is sequentially strict. If the borrower has a bad credit history, it is strictly optimal not to apply, since application is costly and the lender rejects such applications. Given the lender's beliefs, that a borrower with a bad record will default for sure, his strategy is (strictly) sequentially rational.

It only remains to verify that the beliefs of the lender are implied by the D1 criterion. Since $m^* = 1$,³⁰ only this type of borrower will repay. Suppose that the lender lends with probability q to an applicant with signal B , where q is chosen so that an applicant of type $m^* = 1$, who intends to repay, is indifferent between applying or not. Thus, q satisfies

$$(1 - \delta)a = q[1 - \delta + \delta\lambda(V(K) - V(0))]. \quad (\text{A.23})$$

Consider a borrower of type $m' > 1$, whose optimal strategy is to default on the loan, if she receives it. Her net gain from applying, relative to not applying, equals

$$q[(1 - \delta)(1 + g) + \delta(V(K) - V(m' - 1))] - (1 - \delta)a. \quad (\text{A.24})$$

We now show that expression above is strictly positive. Substituting for $(1 - \delta)a$ from (A.23), and dividing by q , we see that the sign of (A.24) is the same as that of

$$(1 - \delta)g + \delta(1 - \lambda)(V(K) - V(m' - 1)) + \delta\lambda(V(0) - V(m' - 1)). \quad (\text{A.25})$$

Since default is optimal for type m' ,

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

³⁰See the proof of Proposition 3.

establishing that the sum of the first two terms in (A.25) is strictly positive. Since $V(0) > V(m')$, the third term is positive, and (A.25) 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 signal. ■

A.7 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

$$\tilde{\pi}(t; K) = \frac{\sum_{m=1}^{m^*(K)} \tilde{\mu}_m(t)}{1 - \tilde{\mu}_0(t)}. \quad (\text{A.26})$$

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

$$\bar{\pi}(K) := \sup_t \tilde{\pi}(t; K) < \frac{\ell}{1 + \ell}. \quad (\text{A.27})$$

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 through 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 $m^\infty := \lim_{K \rightarrow \infty} m^*(K)$, we have that as $K \rightarrow \infty$ 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 $\bar{\pi}(\bar{K}) \rightarrow \frac{1}{\bar{K}}$ as $\lambda \rightarrow 0$.

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.8 Screening borrowers

Fix the credit market interpretation of the trust game, and assume the simple information structure with K period memory, and a pure strategy equilibrium as in Section 4. That is, in equilibrium, all lenders exclude borrowers who have a B signal. We examine the incentives of an individual lender to deviate, by offering a contract to a borrower with a B signal. Our first question is, can such a contract induce separation, i.e. acceptance only by those types m who will repay the loan?

A contract \mathcal{C} with a loan size x specifies a triple $(u(x) - r(x), u(x), v(x))$ denoting borrower utilities in three possible outcomes if the contract is accepted: repayment, voluntary default, and involuntary default. The payoffs to the lender are: $r(x) - (x + f)$ when there is repayment, and $-(x + f)$ when there is default, either voluntary or involuntary, where f is the fixed cost of making a loan.

The borrower has three possible actions: voluntary default, repayment, and refusal of the contract, which we denote by \emptyset . The overall repeated-game payoff to a borrower of type m after these actions are:

$$U(D, m) = (1 - \delta)[(1 - \lambda)u(x) + \lambda v(x)] + \delta V(K),$$

$$U(R, m) = (1 - \lambda)[(1 - \delta)(u(x) - r(x)) + \delta V(m - 1)] + \lambda[(1 - \delta)v(x) + \delta V(K)],$$

$$U(\emptyset, m) = \delta V(m - 1).$$

Thus the payoff difference between repayment and refusal of the loan equals

$$U(R, m) - U(\emptyset, m) = (1 - \delta)[(1 - \lambda)(u(x) - r(x)) + \lambda v(x)] + \delta\lambda[V(K) - V(m - 1)],$$

which is increasing in m .

Similarly, the payoff difference between intended default and refusal, equals

$$U(D, m) - U(\emptyset, m) = (1 - \delta)[(1 - \lambda)u(x) + \lambda v(x)] + \delta[V(K) - V(m - 1)],$$

which is also increasing in m . Thus, any contract that is acceptable to a borrower of type m is (strictly) acceptable to a borrower of type $m' > m$. This implies that there does not exist a contract that induces acceptance from borrowers with lower m values, and rejection from those with higher m values.

Finally, consider the difference

$$U(D, m) - U(R, m) = (1 - \delta)(1 - \lambda)r(x) + \delta(1 - \lambda)[V(K) - V(m - 1)].$$

This is also increasing in m . More importantly, it is increasing in $r(x)$, the repayment required. Thus, by reducing the gross repayment required from the borrower, i.e. by reducing the size of the loan, the lender can induce a higher repayment rate. Thus, smaller loans can induce higher repayment, but this will not be profitable as long as fixed costs f are large enough.

Observe that higher interest rates (greater repayments for a fixed loan size x) also induce more defaults, as in Stiglitz and Weiss (1981).

Finally, let us consider deferred repayments. Let us assume that the deviating lender offers the same terms as other lenders, but defers repayment to the next period with the repayment terms chosen so that the present value of repayments is the same. From the parameters of the trust game in Figure 1a, the repayment required by other lenders is $\frac{g}{(1-\lambda)}$, and consequently, the deviating lender requires a deferred repayment of $\frac{g}{\delta(1-\lambda)}$. We now examine the repayment incentives of a borrower in response to this deferred repayment contract.

Consider a borrower of type $m > 1$. She will be excluded from the credit market in the next period. Consequently, she will repay a loan made today if

$$(1 - \delta)g < \delta^2(1 - \lambda)[V(m - 2) - V(K)].$$

The right hand side reflects the fact that repayment tomorrow affects her continuation value the period after. Since $\delta < 1$, it follows that if type $m - 1$ defaults on a loan with standard terms, then type m will default on a loan with deferred terms. Thus, every type $m > m^* + 1$ will default.

However, for type $m^* + 1$, the comparison is less straightforward. Let us define:

$$\Omega := \frac{(1 - \delta)g}{\delta(1 - \lambda)}.$$

We need to compare Ω with

$$\beta(m^* + 1) := \delta[V(m^* - 1) - V(K)] = V(m^*) - \delta V(K), \quad (\text{A.28})$$

where the second equality follows from $\delta V(m^* - 1) = V(m^*)$. If $\Omega < \beta(m^* + 1)$, then repayment is optimal for type $m^* + 1$; otherwise, default is optimal. We know that

- a borrower of type $m^* + 1$ will default on the standard contract.
- a borrower of type m^* will repay under the standard contract.

This gives us two inequalities:

$$V(m^*) - V(K) < \Omega < V(m^* - 1) - V(K).$$

From (A.28), $\beta(m^* + 1)$ is greater than the left-most term in the above inequality. Since $\delta < 1$, $\beta(m^* + 1)$ is also smaller than the right-most term in the inequality. Thus we are unable to compare $\beta(m^* + 1)$ and Ω , and we conclude that type $m^* + 1$ may or may not switch her behaviour from defaulting to repaying when offered deferred terms.

Now consider a borrower of type $m = 1$. Since she will have a good signal in the next period, her optimal choice, contingent on being able to repay both loans, is to either repay both or to default on both.³¹ It will be optimal to default if:

$$(1 - \delta)g \frac{1 + \delta}{\delta} > \delta(1 - \lambda)[V(0) - V(K)].$$

The incentive to default — the left-hand side above — is more than doubled as compared to the standard contract, while the right-hand side is the same continuation value as under the standard contract. Thus it will be optimal for type $m = 1$ to default as long as K is chosen

³¹Also, if the borrower involuntarily defaults on one loan, it is optimal to voluntarily default on the other.

so that punishments are not too long. In particular, if $K = \bar{K}$, type $m = 1$ will always default.

To summarise, type $m^* + 1$ may or may not switch from defaulting to repaying, while type $m = 1$ switches from repaying to defaulting. The incentives to repay of other types are unaffected. Finally, recall that the distribution of types conditional on a B signal is uniform. We conclude therefore that by deferring repayment the deviating lender either reduces the overall repayment probability at a B signal, or does not affect it.

References

- ALI, S. N., AND D. A. MILLER (2013): “Enforcing cooperation in networked societies,” *Unpublished paper*. [282].
- 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.
- BHASKAR, V., AND I. OBARA (2002): “Belief-Based Equilibria in the Repeated Prisoners’ Dilemma with Private Monitoring,” *Journal of Economic Theory*, 102(1), 40–69.
- BHASKAR, V., AND C. THOMAS (2018): “Community Enforcement of Trust,” Mimeo, University of Texas at Austin. Available at <https://sites.google.com/site/carotho14/research/stable/bounded-memory-29jan18.pdf>.
- 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.
- 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., 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.
- GREIF, A. (1993): “Contract Enforceability and Economic Institutions in Early Trade: The Maghribi Traders’ Coalition,” *American Economic Review*, 83(3), 525–548.
- 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.
- KARLAN, D., M. MOBIUS, T. ROSENBLAT, AND A. SZEIDL (2009): “Trust and social collateral,” *The Quarterly Journal of Economics*, 124(3), 1307–1361.
- KLEIN, B., AND K. B. LEFFLER (1981): “The role of market forces in assuring contractual performance,” *Journal of Political Economy*, 89(4), 615–641.
- KOVBASYUK, S., AND G. SPAGNOLO (2016): “Memory and Markets,” *Available at SSRN 2756540*.
- 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.
- RAHMAN, D. (2012): “But who will monitor the monitor?,” *American Economic Review*, 102(6), 2767–97.
- SEKIGUCHI, T. (1997): “Efficiency in Repeated Prisoner’s Dilemma with Private Monitoring,” *Journal of Economic Theory*, 76(2), 345–361.
- SPERISEN, B. (2018): “Bounded Memory and Incomplete Information,” *Games and Economic Behavior*, *forthcoming*.

- STIGLITZ, J. E., AND A. WEISS (1981): “Credit rationing in markets with imperfect information,” *The American Economic Review*, 71(3), 393–410.
- 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.
- TIROLE, J. (1996): “A theory of collective reputations (with applications to the persistence of corruption and to firm quality),” *The Review of Economic Studies*, 63(1), 1–22.
- WOLITZKY, A. (2012): “Cooperation with network monitoring,” *Review of Economic Studies*, 80(1), 395–427.