The design of credit information systems

V Bhaskar (UT Austin) Caroline Thomas (UT Austin)

September 23, 2016
Preliminary and Incomplete

Abstract

We examine the performance of large credit markets subject to borrower moral hazard when information on the past behavior of borrowers is subject to bounded memory constraints, and agents are subject to small payoff shocks. With perfect information on past behavior, no lending can be sustained, independent of the length of memory. Partial information on credit histories improves outcomes. By pooling recent defaulters with those who defaulted further in the past, such an information structure disciplines lenders, since they are unable to target loans towards defaulters who are unlikely to offend again. Paradoxically, efficiency can be improved by non-monotonic information structures that pool non-defaulters and multiple offenders. Equilibria where defaulters get a loan with positive probability may also improve efficiency, by raising the proportion of bad risks in the pool of defaulters. To summarize, endogenously generated adverse selection mitigates moral hazard.

JEL codes: D82, G20, L14, L15.

Keywords: credit markets, information design, repeated games, random matching.

1 Introduction

In many markets, information on a participant's past behavior is constrained by *bounded memory,* and other agents cannot observe the participant's records more than K periods ago. This is true in the case of credit markets in most countries. In the United States an individual's credit records cannot be used beyond 10 years, while in the UK, the limit is six years. Several other European countries have an even shorter memory for credit histories. A second variant of bounded memory arises from the policies of internet platforms in summarizing the reputational score of their participants. For example, Amazon lists a summary statistic of seller performance over the past 12 months. Finally, one should note that the broader philosophical appeal of bounded memory, that an individual's transgressions in the distant past should not be permanently held against them. This is embodied in the European Court of Justice's determination that individuals have the "right to be forgotten", i.e. they can ask search engines to delete past records pertaining to them. It is also reflected in legislation in 24 US states that mandates that employers cannot ask job applicants to disclose past convictions.

We study the functioning and design of credit information systems that are constrained by bounded memory. Since borrowers have a variety of sources of finance in a modern economy, we study a model with a continuum of infinitely lived borrowers and many lenders, where each borrower-lender pair interacts only once. Lending is efficient and also profitable, provided that the borrower intends to repay the loan; however, the borrower is subject to moral hazard, and may wilfully default on a loan (additionally, there is a small chance of involuntary default). Thus lending can only be supported if default results in future exclusion from credit markets. Nonetheless, future lenders have little interest in punishing a borrower for past transgressions, and will lend if they anticipate repayment. Thus a borrower can be disciplined from wilful default only if this record indicates that he is likely to default in the future. The bounded memory restriction implies that the lender's information regarding a borrower who she is matched with is an information partition on the set of K-length sequences of possible outcomes of the borrower, where the outcome in each period describes whether the borrower got a loan, and her repayment behavior in the event of getting one.

A natural conjecture is that providing maximal information is best, so that the lender's information partition is the finest one. This turns out to be false – perfect information on past behavior of the borrower, in conjunction with bounded memory precludes any lending. In any equilibrium of the game that satisfies a mild requirement, of being robust to small payoff shocks, no borrower will repay, and thus no lender will lend. This leads us to explore partitional information structures that provide the lender partial information on the past K outcomes of the borrower. Since the borrower is aware of the outcomes in his previous interactions, this bestows him with private information. Partial information can indeed do better. If the seller's information partition only discloses whether the borrower has ever defaulted in the last K periods or not, and if K is sufficiently long, then there exists an equilibrium where a defaulting borrower is excluded for K periods. However, K may be longer than is needed to discipline a defaulting borrower; we also need to discipline the lender so that she is not tempted to lend to borrower with a bad credit history. Efficiency can be improved, somewhat paradoxically, by a non-monotonic information partition, where borrowers with multiple defaults are pooled with non-defaulters, since this provides incentives for defaulters to re-offend. It can also be enhanced in an equilibrium where borrowers with bad credit histories are provided loans with some probability, by increasing the representation of "bad risks" in the population with bad credit histories. Since borrowers who have defaulted relatively recently will default again on obtaining a loan, they are over-represented in the pool of borrowers with bad credit histories, thereby disciplining the lenders. We briefly examine non-partitional information structures and show that these may enhance efficiency.

We also investigate alternative extensive form representation of the borrower-lender interaction. If the privately informed borrower is required to move first, e.g. by making a loan application, then the partial information structure induces a signaling game between borrower and lender. There exist equilibria where the length of exclusion for defaulters can be minimized, which are supported by plausible beliefs satisfying standard refinements.

1.1 Related Literature

There is small but growing literature on the role of information in the functioning of credit markets, that includes Pagano and Jappelli (1993), Padilla and Pagano (1997), Elul and Gottardi (2015) and Kovbasyuk and Spagnolo (2016). Most of the literature focuses on adverse selection, and consequently differs from our focus on the moral hazard problem.[1] Padilla and Pagano show that information provision may be excessive. Kovbasyuk and Spagnolo study a lemons market with Markovian types, where the invariant distribution is adverse enough to result in market breakdown, but initial information may permit trade with some borrowers. They find that bounding memory can improve outcomes in this context.

[1]Elul and Gottardi (2015) consider markets with limited moral hazard as well as adverse selection.

Given our focus on borrower moral hazard, our paper relates to repeated games played in a random matching environment, as studied by Kandori (1992), Ellison (1994) and a large set of subsequent papers. Our analysis departs from this literature in many ways. Since we assume a large (continuum) population, this precludes contagion strategies that are usually used in this literature, and instead, we assume that agents have some information on their partner's past interactions (as in Takahashi (2010) or Heller and Mohlin (2015)). Whereas these papers consider a simultaneous move game, almost exclusively the prisoner's dilemma[2], lender-borrower interaction naturally involves a sequential structure, with a delay between the initiation of the loan and the repayment decision. This turns out to make a significant difference to the analysis. In particular, when the seller's information on the borrower's past behavior is partial, an equilibrium of the overall game induces a dynamic game of imperfect information in a given lender-borrower match.

A second feature of our analysis is that we assume that agents in the stage game are subject to small payoff shocks. Motivated by Harsanyi (1973), we focus on purifiable equilibria of the game without payoff shocks, that are limits of equilibria of the perturbed game as the payoff shocks vanish – in our view, this is a mild requirement. Thus this work builds on and complements Bhaskar (1998) and Bhaskar, Mailath, and Morris (2013), both of which demonstrate that the purifiability restriction places limits on the ability of agents to condition their behavior on payoff irrelevant histories.In contrast with this work, the present paper shows that providing partial information on past histories may permit such conditioning.

Given our focus on bounded memory, our work also relates to the influential literature on "money and memory". Kocherlakota (1998) shows that money and unbounded memory play equivalent roles. Wiseman (2015) demonstrates that when memory is bounded, money can do significantly better.

Ekmekci (2011) studies the interaction between a long-run player and a sequence of short run players, where the long run player is subject to moral hazard and her action is imperfectly observed, and some initial uncertainty about the long run player (as in reputation models). With perfect information on past signals, reputation disappears, as in Cripps, Mailath, and Samuelson (2004). Ekmekci (2011) shows that bounded memory allows reputations to persist in the long run. Liu and Skrzypacz (2014) examine a reputational model of buyer-seller interaction where short lived sellers have bounded information on the sellers past decisions. They demonstrate that equilibrium displays a cyclical pattern, whereby the seller builds up reputation before milking it. Sperisen et al. (2016) extends this analysis by considering

[2]But see Deb and González-Díaz (2010).

non-stationary equilibria, and by also requiring that equilibria be robust to payoff shocks.

Our key finding that providing partial information dramatically improves equilibrium outcomes has some resonance with results in the literature on repeated games with imperfect private monitoring. Kandori (1991) provided an early example where private monitoring permits greater efficiency than perfect monitoring. Sugaya and Wolitzky (2014) provide a more general analysis of this question. The mechanism whereby partial information does better is qualitatively different in our context.

Finally, our work relates the recent interest in information design, initiated by Kamenica and Gentzkow (2011), since we explore how informational structures affect the efficiency of equilibrium outcomes. In contrast with this literature, we focus on partitional information structures (but see section 2.3.1). We also focus on the informational requirements for sustaining the most efficient pure strategy equilibrium. Also, while the distribution of types or over states is usually exogenous in the information design context, the induced distributions over types (or histories) in the present paper are endogenous and interact with the information structure.

2 The model

We model the interaction between borrowers and lenders as follows. There is a continuum of borrowers and a continuum of lenders. Time is discrete and the horizon infinite, with borrowers having discount factor $\delta \in (0, 1)$. The discount factor of the lenders is irrelevant for the analysis. In each period, individuals from population 1 (the creditors) are randomly matched with individuals from population 2 (the borrowers),[3] to play the following sequential-move game. First, the creditor chooses between $\{Y, N\}$, i.e. whether or not to extend a loan. If the creditor chooses N, the game is over with payoffs $(0, 0)$. If the creditor chooses Y, then the borrower invests this in a project with uncertain returns. With a small probability λ, the borrower is unable to repay the loan, i.e. is constrained to default, D. With the complementary probability the borrower is able to repay the loan, and must choose whether or not do so, i.e. must choose in the set $\{R, D\}$, where R denotes repayment and D denotes default. If the borrower repays the payoff to the lender is π_l, and that of the borrower is π_b. If the borrower defaults, the lender's payoff is $-\ell$, independent of the reason for default. The borrower's payoff from default depends on the circumstances under which this occur; when

[3]If the measure of the lenders exceeds that of the borrowers, each lender will be matched with more than one borrower – this causes no complications.

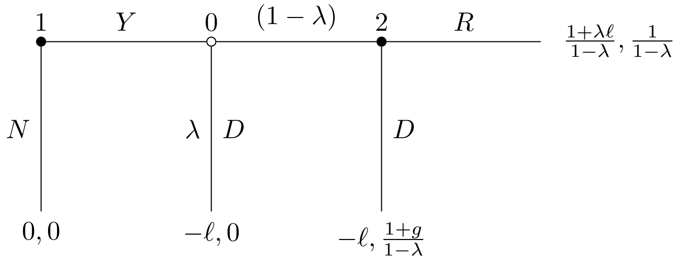

(a) Extensive form.

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

(b) Normal form.

Figure 1: Extensive and normal form representation of the stage-game

he is unable to repay, his payoff is 0, while if he wilfully defaults, his payoff is $\pi_b + \tilde{g}$, where $\tilde{g} > 0$. If the borrower chooses R when he is able to do so, this gives rise to an expected payoff of $(1-\lambda)\pi_b$ for the borrower and $(1-\lambda)\pi_l - \lambda\ell$ for the lender. We normalize these payoffs to $(1, 1)$. The payoffs π_l and π_b are then as in the figure below. If the borrower chooses D when he has a choice, the expected payoff is $(1-\lambda)(\pi_b + \tilde{g})$ for the borrower. Define $g := (1-\lambda)\tilde{g}$, so that the expected payoffs when the borrower wilfully defaults are $(-\ell, 1+g)$. We assume that only the borrower can observe whether or not he is able to repay, i.e. the lender or any outside observer can only observe the outcomes in the set $O = \{N, R, D\}$.

An alternative specification of the model is a follows. The borrower has two choices of project. The safe project results in a medium return, that permits repayment r, with a high probability, $1-\lambda$, and a zero return with probability λ. The risky project results in a high return, $H > M$, but with a lower probability, $1-\theta$, and zero return with complementary probability. Let $(1-\lambda)M - r = 1, (1-\theta)H - r = 1 + g, (\theta - \lambda)r = (1+\ell)$. Thus the expected payoffs are as before, with the information being slightly different (if the borrower chooses the high return project, she still repays with positive probability. We believe that the analysis can be generalized to this case as well, but do not pursue this here.

Since there is a continuum of borrowers and a continuum of lenders, the behavior of any individual agent has negligible effects on the distribution of continuation strategies in the game. Furthermore, since the borrower has a short-term incentive to default, she will do so in any period unless future lenders have information about her behavior. We therefore assume this – the precise details will differ depending on the context we consider.

2.0.1 Payoff shocks: The perturbed game

We now consider a perturbed version of the underlying game. Let Γ denote the extensive form game that is played in each period, and let Γ^∞ denote the extensive form game played in the random matching environment – this will depend on the information structure, which is at yet unspecified. The perturbed stage game, $\Gamma(\varepsilon)$, is defined as follows. Let X denote the set of decision nodes in Γ, and let $\iota(x)$ denote the player who moves, making a choice from a non-singleton set, $A(x)$. At each such decision node $x \in X$ where player $\iota(x)$ has to choose an action $a_k \in A(x)$, his payoff from action a_k is augmented by εz_x^k, where $\varepsilon > 0$. z_x^k is the $k-$th component of z_x, where $z_x \in \mathbb{R}^{|A(x)-1|}$ is the realization of a random variable with bounded support. We assume that the random variables $\{\tilde{z}_x\}_{x \in X}$ are independently distributed, and that their distributions are atomless. Player $\iota(x)$ observes the realization of the shock before he is called upon to move. In the repeated version of the perturbed game, $\Gamma^\infty(\varepsilon)$, we assume that the shocks for any player are independently distributed across periods. [4] In the specific context of the borrower lender game, we may assume that the lender gets a idiosyncratic payoff shock from not lending, while the borrower gets a idiosyncratic payoff shock from repaying. Motivated by Harsanyi (1973), we focus on purifiable equilibria, i.e. equilibria of the game without shocks that are limits of equilibria of the game $\Gamma^\infty(\varepsilon)$ as $\varepsilon \to 0$. In the terminology of Bhaskar, Mailath, and Morris (2013), we require weak purifiability.

Call an equilibrium of the unperturbed game *strict* if a player has strict incentives to play her equilibrium action at every decision node, whether this node arises on or off the equilibrium path.[5] The following lemma is very useful:

Lemma 1 *A strict equilibrium of Γ is purifiable.*

2.1 The infinite memory benchmark

Suppose that each lender can observe the entire history of transactions of the borrower, i.e. if she is matched with the borrower at date t, she observes the outcome of the borrower in periods $1, 2, .., t-1$. Assume that $g < \frac{\delta(1-\lambda)}{1-\delta(1-\lambda)}$. This ensures that there exists an equilibrium where lending takes place, i.e. a borrower who is permanently excluded upon default has an incentive to repay when she is able to.

[4]The assumption that shocks are independently distributed across periods is not essential for the lender, in the context of the lender-borrower game.

[5]This definition is weaker than the usual definition of a strict Nash equilibrium in a strategic form game. If a game has a non-trivial extensive form, and some information set is always unreached under any strategy profile, then no equilibrium can be normal form strict.

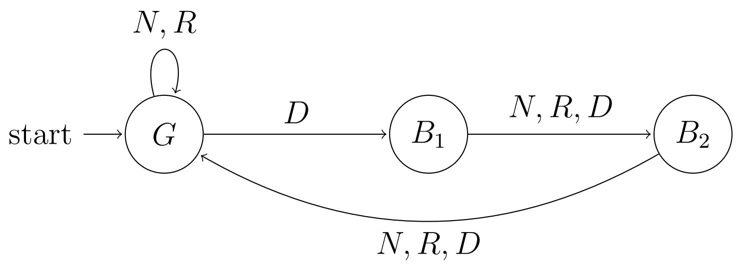

Consider an equilibrium described by the automaton above (for $K = 2$), where a defaulting borrower is excluded for K periods. Depending on the entire history, the borrower is either in a good state or in a bad state, where the bad state is subdivided into K distinct sub-states. The lender extends a loan if an only if he borrower is in the good state. A borrower begins in the good state, and transitions are determined by the outcomes of the stage game: no loan (N), voluntary or involuntary default (D), and repayment (R). A borrower who is in the good state stays there unless she defaults, in which case she transits to the first of the bad states. The borrower then transits through the class of bad states for K periods, independently of the outcome in each of these periods. At the end of K periods, she transits to the good state, again independently of the outcome in the last of these periods. Note that this equilibrium requires that every lender should be able to observe the entire history of every borrower he is matched with, since otherwise, he cannot deduce whether the borrower defaulted in a period where she was supposed to be lent to, or one in which she was supposed to be excluded.

A borrower in a bad state has no incentive to repay, since her continuation value does not depend on whether she repays or not. The incentive constraint, that ensures that she has an incentive to repay in the good state is given by

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

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

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

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

The most efficient equilibrium in this class has K large enough to provide the borrower incentives to repay, but no larger. Call this value $\bar{K}$. Thus 1 is satisfied when $K = \bar{K}$, but is not satisfied with $K = \bar{K} - 1$, implying that

$$(1-\delta)g > \delta(1-\lambda)[V^{\bar{K}-1}(0) - V^{\bar{K}-1}(\bar{K}-1)]. \tag{4}$$

Furthermore, we will also assume, throughout this paper, that 1 holds at as a strict inequality when $K = \bar{K}$ – this assumption will be satisfied for generic values of the parameters (δ, g, λ).

In particular, the most efficient equilibrium payoff in this class equals

$$\bar{V} := V^{\bar{K}}(0) = \frac{(1-\delta)}{1-\delta[\lambda\delta^{\bar{K}} + (1-\lambda)]}. \tag{5}$$

A more efficient equilibrium, with payoff greater than $\bar{V}$, can be sustained if players observe the realization of a public randomization device at the beginning of each period. The best equilibrium payoff for the borrower is one where the incentive constraint just binds. Let V^* denote the best equilibrium payoff a nd let V^{P^*} denote the payoff in the punishment phase. Incentive compatibility requires

$$(1-\delta)g \leq \delta(1-\lambda)(V^* - \hat{V}^{P^*}). \tag{6}$$

Using this, the best equilibrium has payoff

$$V^* = 1 - \frac{\lambda g}{\delta(1-\lambda)}. \tag{7}$$

This payoff can be approximated via a strict equilibrium as follows. A borrower who defaults in a period where the lender should have lent to him is excluded for $\bar{K} - 1$ periods, for sure, and in the $\bar{K}^{\text{th}}$ period with probability $x > x^*$, where x^* is chosen to satisfy

$$V^{P^*} = \delta^{\bar{K}-1}[(1-x^*) + x^*\delta]V^*, \tag{8}$$

$$V^* - V^P = \frac{(1-\delta)g}{(1-\lambda)\delta}. \tag{9}$$

Since $x > x^*$, the equilibrium is strict. Thus we can achieve a payoff arbitrarily close to V^* via an equilibrium that is strict and therefore purifiable. Note that this equilibrium requires that every lender should be able to observe the entire record on the realization of the public randomization device, since otherwise, he cannot deduce whether the borrower defaulted in a period where he was supposed to be lent to, or one in which he was supposed to be excluded. [6]

Bounded Memory Henceforth, we shall assume that lenders have bounded memory, i.e. we assume that the lender observes a bounded history, of length K, of past play of the borrower. We begin by considering partitional information structures, so that, in each period, the lender observes a partition of O^K, where this observation pertains to the past K outcomes of the borrower he is currently matched with. We assume that the lender does not observe any information regarding other lenders. Specifically, he does not observe any information regarding the lenders that the borrower he currently faces has been matched with in the past.

We now show that payoff shocks imply that the borrower's information regarding the current lender is irrelevant, and cannot be conditioned upon in any purifiable equilibrium. Let the information about the lender that is observed by the borrower that she is matched within a period be arbitrary, and denote it by h_1. Suppose that the lender lends to the borrower. Since no future lender observes h_1, the borrower's continuation value does not depend upon h_1. Thus in the perturbed game, the borrower can play differently after two different lender histories, h_1 and h_1', only for a set of payoff shocks that have Lebesgue measure zero. Hence, in any purifiable equilibrium of the unperturbed game, no borrower conditions on any information regarding the lender. So we may as well assume that the borrower observes no information regarding the lender.

2.2 Perfect Bounded Memory

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

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

The proof is an adaptation of the argument in Bhaskar, Mailath, and Morris (2013). Suppose that the information partition is the finest possible. When $K = 1$, then the borrower

[6] The payoff V^* can be achieved via a mixed equilibrium that does not require a public randomization device. If the borrower is always indifferent between defaulting and repaying, then he is willing to repay when the loan should have been made, and default when it should not have been made. However, with shocks to his stage game payoff from default, he would condition on shocks and not on whether the loan should have been made. Thus, such an equilibrium is not purifiable.

who receives a loan in period t knows that the lender at $t+1$ cannot observe the outcome at $t-1$. The payoff shocks imply that for any strategy of the lender at $t+1$, the borrower is indifferent between R and D on a set of measure zero, and thus cannot condition on the outcome at $t-1$ in any optimal strategy. This implies that the lender at date t can only be indifferent between Y and N on a set of measure zero, and thus cannot also condition upon the outcome at date $t-1$. The argument for the case when K is arbitrary generalizes this idea, using induction. When $K=1$, the first step in the above argument applies to any partition of O.

A second implication of the proposition is that we need $K \geq 2$, since otherwise no information structure sustains lending. So even if $g < \frac{\delta(1-\lambda)}{1+\delta\lambda}$ so that only one period of memory is required to satisfy (1), we need at least two period memory.

Surprisingly, *less information* may support cooperative outcomes, as we shall see below. Although it still remains the case that the borrower will not condition on events that happened exactly K periods ago, a coarser information structure prevents the lender from knowing this, and thus the induction argument underlying the above proposition (and the main theorem in Bhaskar, Mailath, and Morris (2013)) does not apply.

2.3 Imperfect Information enhances efficiency

We are interested in investigating equilibria where lending is sustained, and where a borrower who defaults involuntarily is not permanently excluded from credit markets, i.e. denial of credit is only temporary. This concern is motivated by the possibility of involuntary default – even if a borrower intends to repay, with probability λ, she will be unable to do so. We therefore investigate the possibility of equilibria with temporary exclusion.

Let K denote the length of memory, and assume that $K \geq \max\{\bar{K}, 2\}$. Suppose that the information structure is a binary partition of O^K, and the creditor observes a "bad" signal if and only if the borrower has had an outcome of D in the last K periods, and observes a "good" signal otherwise. Let B and G denote the respective information sets. This information structure is *monotone*, since if bounded histories are ordered by default rates, then we do not pool worse types with good types, while separating intermediate types.

The borrower has complete knowledge of her own private history. We partition the set of private histories, into $K+1$ equivalence classes, $(I_m)_{m=0}^{K}$. Let t' denote the date of the most recent incidence of D in the borrower's history, and let $j = t - t'$, where t denotes the current period. Define $m := \min\{K+1-j, 0\}$. Thus the borrower knows that if $m = 0$, the lender observes G, while if $m \geq 1$, the borrower observes B.

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 $m \geq 1$, since a borrower with private history I_m is excluded for m periods before getting a clean history, we have that

$$V^K(m) = \delta^m V^K(0), \qquad m \in \{1, \ldots, K\}. \tag{10}$$

Since $K \geq \bar{K}$ (cf. equation 1), the borrower's incentive constraint is satisfied with strict inequality at a good credit history. We examine the borrower's repayment incentives when the lender sees a bad credit history. Note that this is an unreached information set at the candidate strategy profile, since the lender is making a loan when he should not. Repayment incentives depend upon the borrower's private information, and are summarized by m. Note that the borrower's incentives at $m = 1$ are identical to those at $m = 0$ – if $m = 1$ and the borrower repays, then the credit history in the next period will be G. Thus borrowers of type $m = 1$ will always choose R. Now consider the incentives of a borrower of type $m = K$. We need this borrower to default (since otherwise, every type of borrower would repay and then lending after B would be optimal). Thus we require

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

or

$$g \geq \delta^K(1-\lambda)V^K(0). \tag{11}$$

It can be shown that (11) is satisfied when $K = \bar{K}$. (Indeed, (11) is implied by (1) and (4)). Since $V^K(0)$ is decreasing in K, (11) is therefore satisfied for every $K \geq \bar{K}$.

Consider the incentives to repay on being given a loan, at an arbitrary m, when the borrower is able to repay. Choosing D over R entails the payoff difference

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

When $m = K$, the above expression is positive, since we have assumed that $K \geq \bar{K}$, so that (11) is satisfied and the borrower has an incentive to default at K. When $m = 1$, the expression is negative since a borrower has strict incentives to repay after a good credit history. Thus there exists a real number, $m^*(K) \in (1, K)$ that sets the payoff difference to zero:

$$m^*(K) = \ln\left[\frac{(1-\delta)g}{(1-\lambda)V^K(0)} + \delta^{K+1}\right] / \ln\delta. \tag{12}$$

We assume that $m^*(K)$ is not an integer, as will be the case for generic payoffs.[7] If $m < m^*(K)$, the borrower chooses R when offered a loan. If $m > m^*(K)$, she chooses D.

Now consider the beliefs of the lender when she observes B. In each period, the probability of involuntary default is constant, and equals λ. Furthermore, under the candidate profile, a borrower who defaults never gets a loan and hence transits deterministically through the states $m = K, K-1, .., 1$. Thus in any time period $t \geq K$, the lender assigns equal probability to the borrower having m taking the values $1, \ldots, K$.[8] Conditional on B, he attributes probability $\lfloor m^*(K) \rfloor / K$ to the borrower choosing R if she is able to repay, where $\lfloor x \rfloor$ denotes the largest integer no greater than x. His expected payoff from making a loan is therefore

$$\Pi(K) := \frac{\lfloor m^*(K) \rfloor}{K} - \left(1 - \frac{\lfloor m^*(K) \rfloor}{K}\right)\ell, \tag{13}$$

Thus the candidate equilibrium exists as long as $\frac{\lfloor m^*(K) \rfloor}{K}$ is sufficiently small so that the above payoff is strictly negative. Notice also that in periods $t < K$, the lower values of K are censored, and thus if $\Pi(K) < 0$ in (13), it is also not profitable to lend to a borrower with a bad credit history in the initial periods of the game.

We now consider the comparative statics of m^* with respect to K. First, for every (δ, λ, g) such that $g < \frac{\delta(1-\lambda)}{1+\delta\lambda}$, we have that $1 \leq m^*(\bar{K}) < 2$, so that $\lfloor m^*(\bar{K}) \rfloor = 1$. Second, note that $m^*(K)$ increases more slowly than K, since

$$\frac{dm^*}{dK} = \frac{\delta^{K+1}}{\delta^{m^*}}\left(1 - \frac{g\lambda}{1-\lambda}\right) < 1.$$

The last inequality follows from the fact that $m^* < K+1$ and $g < \frac{\delta(1-\lambda)}{1+\delta\lambda}$. Furthermore, since

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

we have that

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

Therefore, $\frac{\lfloor m^*(K) \rfloor}{K} \to 0$ as $K \to \infty$. We conclude that lending is never profitable, for suffi-

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

[8] In fact, under the pure strategy profile described, the invariant distribution over {0,1,2,..,K} has $\mu_0 = \frac{1}{1+K\lambda}$ and $\mu_1 = \cdots = \mu_K = \frac{\lambda}{1+K\lambda}$.

ciently large K. Now for generic values of δ and g, $m^*(K)$ is not an integer, and thus the borrower has strict incentives at each information set. Also, the lender has strictly positive payoffs from making a loan after a good credit history, and strict losses from making a loan after a bad credit history. Thus the equilibrium we have constructed in the game without payoff shocks is strict – an agent has strict incentives to comply with his or her equilibrium action at each information set, regardless of whether this is reached or not. Thus, if the payoff shocks are parameterized by ε as sufficiently small, the equilibrium in the perturbed game induces the same behavior. Thus the equilibrium is purifiable. We therefore conclude:

Proposition 3 *An equilibrium where the lender lends after observing G and never lends on observing B always exists as long as K is sufficiently large. Such an equilibrium is strict and thus purifiable.*

Recall that the maximally efficient pure strategy equilibrium requires $\bar{K}$ periods of punishment, while the equilibrium constructed here may require $K \geq \bar{K}$ periods of punishment. To get an idea of how the inefficiency depends upon the parameters, the following figure plots $\frac{\lfloor m^*(K) \rfloor}{K}$ as a function of K for different values of g. The graph in each case only begins at $\bar{K}$. We see that for small values of g, $\bar{K}$ is small, and $\frac{\lfloor m^*(K) \rfloor}{K}$ is large, so that lending will be profitable for K values near $\bar{K}$. For large values of g, $\bar{K}$ is also large, and $\frac{\lfloor m^*(K) \rfloor}{K}$ is small at $\bar{K}$. Thus lending will be unprofitable at $\bar{K}$, and the equilibrium is likely to be close to the maximally efficient pure strategy equilibrium.

2.3.1 Non-partitional information

We have restricted attention to partitional information structures on O^K. Non-partitional information structures could possibly do better, in terms of the efficiency of outcomes. Suppose there exists a pure strategy equilibrium with $\bar{K}$ period punishments. Consider the following signal structure that only depends on $\bar{K}$ period histories. If there is no instance of D in the last $\bar{K}$ periods, signal G is observed by the lender. If there is any instance of D in the last $\bar{K} - 1$ periods, then signal B is observed. Finally, if there is a single instance of D in the last $\bar{K}$ periods and this occurred exactly $\bar{K}$ periods ago, signal G is observed with probability $(1 - x)$, and B is observed with probability x.

Consider the pure strategy profile where the lender lends if and only if he observes G, and where the borrower chooses to repay if and only if the lender observed G. In this case we have

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

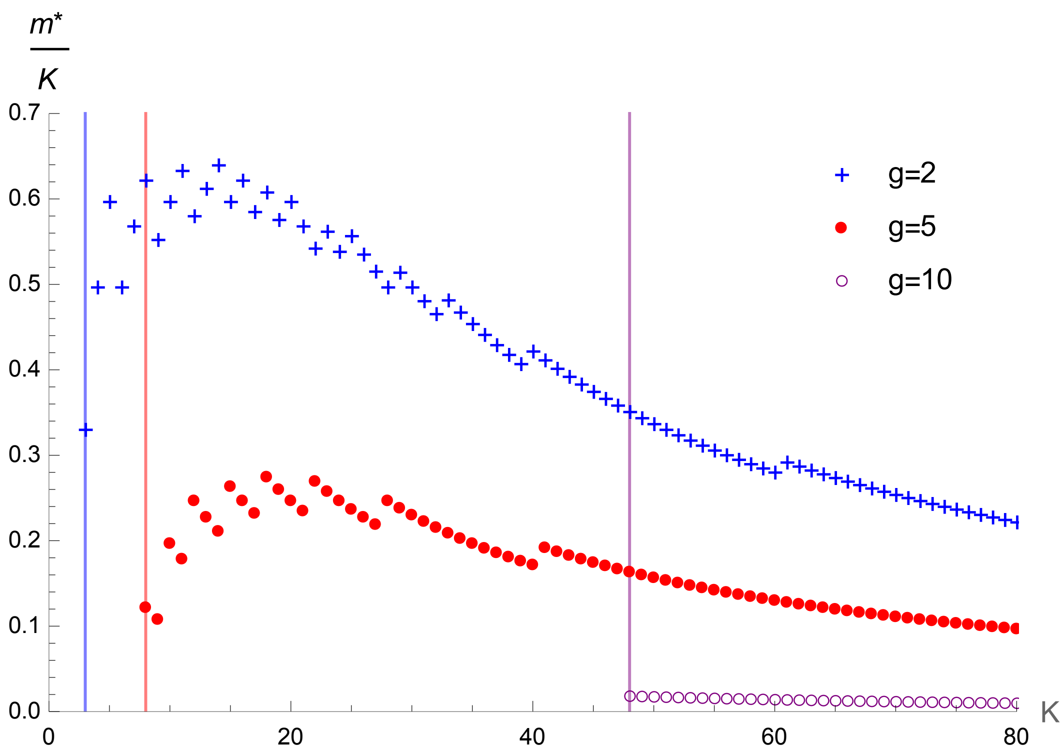

Figure 2: *Proportion* $\frac{\lfloor m^*(K) \rfloor}{K}$ *of non-defaulters among borrowers with bad credit history, for* $g = 2$, $g = 5$ *and* $g = 10$. *The vertical lines represent* $\bar{K}$ *in each case. (Illustrated for* $\delta = 0.91$.*)*

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

so that

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

The borrower agrees to repay at G if and only if

$$(1-\delta)g \leq \delta(1-\lambda)\left(1-\delta^{\bar{K}-1}(\delta x+1-x)\right)\ V^{\bar{K},x}(0). \tag{14}$$

The right hand side of (14) is a strictly increasing function of x. Moreover, by the definition of $\bar{K}$, (14) is satisfied when $x = 1$, and violated when $x = 0$. There therefore exists $x^*(\bar{K}) \in (0, 1]$ such that (14) is satisfied for every $x \in [x^*(\bar{K}), 1]$. (In the non-generic cases where $\bar{K}$ satisfies (1) with equality, we have that $x^*(\bar{K}) = 1$.)

It remains to examine whether this modification preserves the incentive of lenders not to lend at B. While $m^*(\bar{K}, x)$ decreases with x, it remains an element of the interval $[1, 2)$, so that $\lfloor m^*(\bar{K}, x)\rfloor = 1$ 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 conclude therefore that a non-partitional information structure will do better. In particular, whenever there is a $\bar{K}$ memory partitional information structure that works, the non-partitional signal structure set out here can approximate V^*, the efficient borrower payoff.

One concern with non-partitional information structures is that they may be vulnerable to manipulation. Since the report in the final period of the punishment is random, influential borrowers maybe able to persuade the credit agency (information provider) to modify the report in their favour. This would still be consistent with the provider having a fraction x of period K reports being B. If $K > \bar{K}$, this manipulation does not destroy the incentives for repayment by such influential borrowers. However, when the punishment is of length $\bar{K}$, an influential borrower's incentive constraint would be violated, since her effective punishment is of length $\bar{K} - 1$.

2.4 Non-monotone information structures

Recall that $\bar{K}$ periods of exclusion are sufficient to provide incentives for the borrower not to wilfully default. Nonetheless, the equilibria constructed in the previous section may well need $K > \bar{K}$. These are due to two reasons, but both are ultimately tied to ensuring that lenders do not lend to borrowers with a recent history of default. First, we need to ensure that at least some borrowers with a bad credit history do not repay, and thus memory needs to be longer than $\bar{K}$. Second, we need to ensure that sufficiently many types of borrower among those with a bad credit history do not repay so that $\Pi(K)$ is negative. We now show that both these constraints can be considerably relaxed, and in some cases, eliminated, if we permit information structures that are both non-monotone and non-partitional – non-monotonicity plays a significantly more important role in this argument. The idea is as follows. Suppose that we reward borrowers for defaulting on a lender who should not have made a loan. This makes it more likely that a borrower with a bad credit history will default, and makes lending to them less attractive.

Consider the following binary partition of K period histories, where $K \geq \max\{\bar{K}, 2\}$. Let N_D denote the number of instances of D in the last K periods. The credit history is G if $N_D \in \{0, 2\}$, and is B otherwise. The equilibrium strategies are as follows. The lender lends on seeing G and does not after B. Thus if $N_D = 0$, a borrower has every incentive to repay since if he plays D, he will be excluded for K periods. Consider now the repayment incentives of a borrower with $N_D = 1$. His incentives depend on $m \in \{1, 2, .., K\}$, where m is defined as before, as the number of periods remaining before the single default disappears from the record, conditional on not defaulting again. Now if $m = 1$, the record of the past default will be wiped clean tomorrow if he repays, but will be renewed if he defaults since this results in a next period signal B. Thus the borrower will repay if $m = 1$. Now consider $m > 1$. If the borrower defaults, then in the next period, there will be two instances of D, and thus a good credit history, which will remain as long as he does not default, for $m - 1$ periods. Furthermore, it can be renewed at this point by defaulting! In the appendix we show that his optimal continuation strategy is somewhat complex, it can be shown that defaulting today raises the borrower's current payoff as well as his continuation value, whenever $m > 1$. Consequently, the expected payoff of the lender from making a loan to a borrower with credit history B, under the invariant distribution induced by the strategy profile is

$$\frac{1}{K} - \frac{K-1}{K}\ell,$$

which must be less than zero for such lending to be unprofitable. If this is negative when $K = \bar{K}$, then there exists an equilibrium that achieves the best pure strategy equilibrium payoff, $\hat{V}$. Otherwise, the best equilibrium in this class is the smallest value of K that makes profits in the above expression negative.

This result can be strengthened, by allowing non-partitional information structures. Consider the following implementation of the minimal punishment that satisfies incentive compatability, where after a default, the borrower is excluded for $\bar{K}-1$ periods for sure, is excluded with probability x^*. Recall from section 2.1 that x^* is chosen to solve equations (8) and (9).

Consider the following signal structure observed by the lender, regarding the borrower's history. Signal G is sent if $N_D \in \{0, 2\}$. Signal B is sent if $N_D > 2$, or $N_D = 1$ and $m > 1$. If $N_D = 1$ and $m = 1$ so that the single instance of D occurred $\bar{K}$ periods ago, then signal G is sent with probability $1 - x$ and signal B is sent with probability $x > x^*$. Our preceding arguments imply that if $N_D = 0$, the borrower has strict incentives to to repay, since $x > x^*$. The borrower with bad signal and $N_D = 1$ defaults unless the single default happened exactly $\bar{K}$ periods ago. By choosing x arbitrarily close to x^*, the lender's payoff from lending to a borrower with a bad signal can be made arbitrarily close to

$$\frac{x^*}{\bar{K}} - \frac{\bar{K}-1}{\bar{K}}\ell, \tag{15}$$

If the above expression (15) is strictly negative, we have an equilibrium that can approximate the efficient payoff V^* arbitrarily closely.

If this is not the case, and the expression in (15) is positive, we cannot approximate V^*. However, we can approximate the best pure strategy equilibrium payoff, $\hat{V}$, where punishments are essentially of length $\bar{K}$. Assume that memory is of length $\bar{K}+1$. In the case of exactly one instance of D, signal B is sent with probability one if this instance is in the last $\bar{K}$ periods and with probability x if the instance is exactly $\bar{K}+1$ periods ago. Let x be very close to zero. Then the profits for the lender from lending to an agent with a bad signal is

$$\frac{x}{\bar{K}+1} - \frac{\bar{K}}{\bar{K}+1}\ell < 0.$$

Thus a strict equilibrium exists, that achieves a payoff $\hat{V}$, arbitrarily close to the pure strategy equilibrium. We summarize our results in the following proposition.

Proposition 4 *If $\bar{K} \geq 2$, there exists binary non-monotone non-partitional information structures, and associated strict equilibria that can approximate $\hat{V}$, the best pure strategy purifiable equilibrium. If x^* is small, or $\bar{K}$ and ℓ are large so that $x^* - (\bar{K}-1)\ell < 0$, there exists a non-monotone signal structure that can approximate the fully efficient payoff, V^*.*

Non-monotonicity of the information is perhaps an unattractive feature, since a borrower with a worse default record is (implicitly) given a better rating than one with a single default. Furthermore, it may be vulnerable to manipulation, if we take into account considerations that are not explicitly modelled. A borrower with two defaults in the last $K-1$ periods is ensured of a continuation value corresponding to being able to default without consequence every K periods. Thus for the rest of the paper we restrict attention to monotone information structures.

2.5 A mixed equilibrium when bad applicants are too profitable

We now return to the case of a monotone information structure, where credit history B is observed if there is any default in the last K periods, with G being observed otherwise. Let us assume $K \geq \bar{K}$, so that memory is sufficiently long for the borrower to be disciplined. However, there may not be an equilibrium where a borrower with a bad credit history never receives a loan, since the lender may find lending too profitable. Longer memory allows us to provide incentives to the lender, but at the cost of inefficiency, since the punishments for defaulters are longer than necessary. Therefore, we consider an equilibrium where the lender lends with probability 1 on observing G and probability $p \in (0,1)$ on observing B. If p is small, then punishments are severe enough that a borrower with a good credit history has incentives to repay, as does a borrower with bad credit history who has defaulted long ago, i.e. when m is small. On the other hand, borrowers who have defaulted relatively recently will default upon getting a loan. Since $p > 0$, this "tilts" the invariant distribution over $m \in \{1, 2.., K\}$ so that higher values of m have greater probability, thereby disciplining the lenders. Since we have to explicitly deal with the exogenous default probability λ in computing the invariant distribution, this is reflected in our analysis below.

Let $p \in [0,1]$ be the probability that a lender with a bad credit history gets a loan (the lender with a good credit history gets a loan for sure). Let $\bar{m} \in \{0, 1, 2, .., K+1\}$ be the smallest value of m such that the borrower finds it optimal to default. Now all but a finite number of values of p, the borrower will not be indifferent between repaying and defaulting at $\bar{m}$. In this case, the borrower's strategy is to repay if $m < \bar{m}$ and to default if $m \geq \bar{m}$. Thus the value function is given by the recursion

$$V(m) = p(1-\delta) + \delta\left[p\lambda V(K) + (1-p\lambda)V(m-1)\right] \qquad \text{if } m < \bar{m},$$

$$V(m) = p(1-\delta)(1+g) + \delta\left[pV(K) + (1-p)V(m-1)\right] \qquad \text{if } m \geq \bar{m}.$$

In equilibrium, $\bar{m}$ is the smallest integer in $\{0, 1, 2, .., K + 1\}$ for which the borrower weakly prefers D to R at $\bar{m}$, conditional on obtaining a loan and being able to repay, so that

$$(1-\delta)g \geq \delta(1-\lambda)\left[V(\bar{m}-1) - V(K)\right].$$

However, for non-generic values of p, the relation above holds with equality and the borrower is indifferent between D and R at $\bar{m}$. In this case, the borrower's set of best responses are summarized by $(\bar{m}, q)$, $q \in [0, 1]$. That is the borrower can default at $\bar{m}$ with probability q, while defaulting for sure at $m \geq \bar{m}+1$ and repaying for sure at $m < \bar{m}$. Notice that the best responses of the borrower only depend on p. Let $\rho(p) \in \{0, 1, 2, \ldots, K+1\} \times [0, 1]$ denote a best response of the borrower.

Given the exogenous default probability λ, the pair $(p, \rho(p))$ induces an invariant distribution μ on the state space $\{0, 1, 2, \ldots, K\}$. For $0 < j < \bar{m}$, a lender transits to $j-1$ unless he gets a loan and suffers involuntary default, i.e.

$$\mu_{j-1} = (1-p\lambda)\mu_j \text{ if } j < \bar{m}.$$

For $j > \bar{m}$, a borrower only transits to $j-1$ if she does not get a loan, so that

$$\mu_{j-1} = (1-p)\mu_j \text{ if } j > \bar{m}.$$

If $j = \bar{m}$,

$$\mu_{\bar{m}-1} = [(1-p) + p(1-\lambda)(1-q)]\mu_{\bar{m}} \ .$$

The measure of μ_K equals both the inflow of involuntary defaulters, who defect at rate λ and the inflow of deliberate defaulters from states $j \geq m$, and equals

$$\mu_K = \lambda\left(\mu_0 + p\sum_{k=1}^{K}\mu_k\right) + p(1-\lambda)\sum_{j=\bar{m}+1}^{K}\mu_j + p(1-\lambda)q\mu_{\bar{m}}.$$

Finally, the measure of agents with a good credit history equals

$$\mu_0 = (1-\lambda)\mu_0 + (1-p\lambda)\mu_1.$$

Let $\mu(p, \bar{m}, q)$ denote the invariant distribution as a function of p and a best response $\rho(p) = (\bar{m}, q)$. Let $\hat{\mu}(p)$ denote the set of invariant distributions that arise from some $\rho(p)$ – this correspondence is upper-hemicontinuous and convex valued.

Conditional on being given a loan, the probability that a borrower with credit history B defaults is given by

$$\pi(p;\bar{m},q;\mu) = \frac{\sum_{j=\bar{m}+1}^{K} \mu_j + q\mu_{\bar{m}}}{1-\mu_0}.$$

A mixed equilibrium exists if there is a p and associated $\rho(p) = (\bar{m}, q)$ and $\hat{\mu}(p)$ such that

$$[1 - \pi(p;\rho(p);\hat{\mu}(p))] - \pi(p;\rho(p),\hat{\mu}(p))\ \ell = 0.$$

If the expression on the left hand side is greater (resp. less) than zero, then the lender's best response is to choose $p = 1$ (resp. 0).

Proposition 5 *Assume that $K \geq \bar{K}$, and that there is no pure strategy equilibrium where a borrower who defaults is excluded for K periods. There exists a mixed equilibrium, where a borrower whose credit history indicates default in the last K periods is excluded with probability $p < 1$. Furthermore the borrower who has not defaulted in the last K periods has strict incentives to repay.*

Proof. See Appendix. ■

The above equilibrium is not strict, being in mixed strategies. We conjecture (this remains to be proven) that the equilibrium is nonetheless purifiable, since no agent ever conditions on payoff irrelevant information.

Comparing the mixed equilibrium with $\bar{K}$ period with a pure equilibrium which requires a strictly larger K, we see that the borrowers are strictly better off and has a strictly higher payoff in this equilibrium, whenever $\lambda > 0$. Such a borrower could always follow a strategy of never defaulting, when he has a bad credit history. By doing so, he is excluded only partially (since p is interior) and also for a shorter period of time. For lenders, since they are indifferent between making and not making loans to borrowers with a bad credit history, the equilibrium with a larger proportion of borrowers with a good credit history, μ_0, is the one that yields higher welfare. Here the effects are ambiguous – although K needs to be longer for a pure strategy equilibrium, the fact that borrowers may default again when they have a bad credit history makes for greater persistence.

2.5.1 Illustration

Suppose $\delta = 0.9$, $\lambda = 0.1$, $g = 2$ and $\ell = 0.315$. In this case we have that $\bar{K} = 4$.

Consider the pure strategy profile at which $K = \bar{K}$, and at which the lender extends a loan if and only if he observes G. The invariant distribution over $\{0, \dots, K\}$ is given by

$$\mu_0 = 0.714286, \quad \mu_1 = \cdots = \mu_4 = 0.0714286.$$

Under this strategy profile, the lender's expected payoff is equal to the probability of encountering a borrower with a clear history: $\mu_0 = 0.714286$. The expected payoff to a borrower with a clear history is $V^K(0) = 0.763644$.

However, this strategy profile is not an equilibrium. If a borrower with a bad history *were* extended a loan, and was able to repay, she would do so whenever $m \geq 2$, since $m^*(K) = 1.19721$. If the borrower deviated, and extended a loan also at B, his expected payoff would therefore be

$$\Pi(K) = \frac{1 - 3\ell}{4} = 0.01375 > 0.$$

This deviation is therefore strictly profitable.

There exists a mixed strategy equilibrium with $K = \bar{K}$ at which the lender offers a loan with probability $p = 0.033501632$ to a borrower upon observing B (and with probability 1 upon observing G). In this equilibrium, $\bar{m} = 2$, and the invariant distribution over $\{0, \dots, \bar{K}\}$ is given by

$$\mu_0' = 0.704909, \quad \mu_1' = 0.0706872, \quad \mu_2' = 0.0727055, \quad \mu_3' = 0.0747815, \quad \mu_4' = 0.0769167.$$

At this equilibrium, the expected payoff to the lender equals the probability of observing a good signal: $\mu_0' = 0.704909$. This payoff is lower than his expected payoff under the pure strategy profile. Conversely, the expected payoff to a borrower who has $m = 0$ is $V(0) = 0.775268$. Because the borrower is sometimes extended a loan even when the signal is B in the mixed strategy equilibrium, her payoff cannot be lower than when she is never extended a loan at a bad signal.

A pure strategy equilibrium exists if K is sufficiently large. The smallest value of K at which a pure strategy equilibrium exists is $K = 30$. In this case, $m^* = 7.53891$ and the lender strictly prefers not lending at B:

$$\Pi(K) = \frac{7 - 25\ell}{30} < 0.$$

The invariant distribution over $\{0, \ldots, K\}$ is given by

$$\mu_0'' = 0.25, \quad \mu_1'' = \cdots = \mu_4'' = 0.025.$$

The lender's expected payoff is therefore $\mu_0'' = 0.25$, and the expected payoff of a borrower with a clear history is $V^K(0) = 0.537101$. Because she is never offered a loan at B, and because the punishment phase lasts longer, the borrower's expected payoff is much lower than under the mixed strategy profile.

2.6 The length of memory

Suppose that the bound on information, K, is exogenously given. How would we modify the information partition in order have an effective memory $K' < K$? This is relevant, for example, when $K > \bar{K}$, and we seek to implement an equilibrium with $\bar{K}$ period memory. One way of doing this is obvious – construct an information partition (such as the binary partition in our previous sections) where the partitioning does not depend upon the last $K' - K$ periods. A second way is to release *complete information* on the outcomes in periods $K+1, .., K'$. That is, we construct an information partition with $2 \times |O|^{K'-K}$ elements, which indicates whether a default took place in the most recent K periods or not, and reveals the exact outcome in each period after that. The induction arguments that underlie proposition 1 can be used to show that in any equilibrium, there cannot be any conditioning on the outcomes in periods $K + 1, .., K'$.

Suppose that $K < \bar{K}$. In this case, we cannot have an equilibrium with K periods of exclusion, since this is inadequate to provide incentives. The following equilibrium ensures incentives by providing infinitely lived exclusion, as long as $K \geq 2$. Partition the set of outcomes in O^K into those where there is R in every period, and the complement (where there is either N or D in one of the last K periods). The strategy of each lender is to lend to borrower with "always R", and not lend when she observes the complement.

We verify that the borrower who has always played R will play R on getting a loan. His expected discounted payoff is approximately 1 since λ is close to zero, while if he defaults, is payoff is $(1 - \delta)(1 + g)$, since default results in permanent exclusion, which is strictly lower since $g < \frac{\delta}{1-\delta}$.

Consider now a borrower, the credit history about whom is "not always R". Suppose that the borrower's outcome is not R in at least one of the last $K - 1$ periods. Then regardless of the borrower's current decision, the credit history regarding his history will be "not always

R" in the next period. Thus his continuation value is 0 independent of his current action and it is optimal to play D. The only information set where a borrower whose credit history is "not always R" wants to repay is one where the outcome is not R exactly K periods ago, and has been R ever since. But such a history has zero probability under the strategy profile, even when $\lambda > 0$, since $K \geq 2$. Thus the lender assigns probability one to default by a borrower whose credit history is "not always R", and will never lend.

3 An alternative specification of lender-borrower interaction

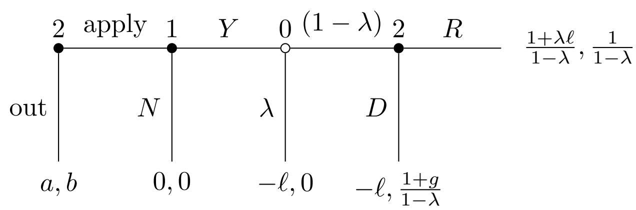

Our modelling of the stage game played between lender and borrower is a simplified version of what usually happens in reality. Usually, a borrower has to make a loan application before this is evaluated by the lender. It turns out that this realistic variation in the extensive form makes a difference. We model the game between borrower and lender as follows. First, the borrower has to apply for a loan, at a small cost. If she does not apply, the game ends, with payoffs $a > 0$ for the borrower and b for the lender. Once she applies, the game is as before; i.e. the lender chooses between Y and N, and if the borrower is able to repay, he must choose between R and N. Thus the backward induction outcome is OUT, where the borrower does not apply. Assume the information structure on the borrower where the lender is informed only whether the borrower has defaulted in the last K periods (credit history G) or not (credit history B).

This modelling transforms the interaction between an individual lender and borrower to a signaling game. In particular, a borrower with credit history B has private information regarding m, the number of periods of no-default that must expire before his credit history becomes good. Thus, in the context of an equilibrium where buyers with different credit histories are treated differently, the application decision can signal the buyer's private information.

Proposition 6 *If $K \geq \max\{\bar{K}, 2\}$, there exists a pure strategy perfect Bayesian equilibrium* *where borrowers with a good credit history apply, while those with a bad credit history do not. The lender lends to an applicant with a good credit history and does not lend to one with a bad credit history. Furthermore, beliefs satisfying the D1 refinement imply that the lender assigns probability one to an applicant with a bad credit history defaulting.*

Proof. The borrower's strategy is a strict best response to the lender's strategy, since application is costly given $a > 0$. Given the lender's beliefs, his strategy is (strictly) sequentially rational. It remains to verify that the beliefs of the lender are implied by the D1 criterion. Suppose that borrower with a bad credit history of type m gets a loan after applying. As in section 2.3, his repayment incentives are such that he will repay if $m < m^*(K)$ and default if $m > m^*(K)$. Consider a mixed response of the lender to a loan application, whereby he gives a loan with probability q on observing a B applicant, where q is chosen so that some type of applicant who intends to repay (i.e. one the types with $m < m^*$) is indifferent between applying or not. Thus, q satisfies

$$a = q[(1 - \delta) + \delta\lambda(V(K) - V(m - 1))], \tag{16}$$

Now consider a borrower of type $m' > m^*$, whose optimal strategy is to default on the loan, if she receives it. Since default is optimal, his net benefit from applying equals

$$q[(1 - \delta)(1 + g) + \delta[V(K) - V(m' - 1)] - a. \tag{17}$$

We show that expression (17) above is strictly positive. Substituting for a from 16, we see that the sign of 17 is the same as that of

$$(1 - \delta)g + \delta[(1 - \lambda)(V(K) - V(m' - 1))] + \delta\lambda(V(m - 1) - V(m' - 1))]. \tag{18}$$

Since default is optimal for type m',

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

establishing that the sum of the first two terms in 18 is strictly positive. Also, $V(m) > V(m')$ since $m' > m$, and so the third term and the overall expression in 18 is strictly positive. So 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 (Banks and Sobel (1987)) implies that in the above equilibrium, the lender must assign probability one to defaulting types when she sees a B credit history applicant. ■

The above proposition implies that a pure strategy equilibrium with $\bar{K}$ periods of exclusion – the minimal number required to provide incentives for repayment – always exists as long as $\bar{K} \geq 2$. If the costs of application a and b are small, and if parameter values are such that only a mixed strategy equilibrium exists at K, as in sub-section 2.5, then a similar mixed strategy equilibrium also exists here. In such an equilibrium, all applicants apply, and an applicant with a bad credit history is given a loan with probability p, just as in sub-section 2.5.

4 A General Class of Two-Player Games

We show in this section that our arguments generalize to a class of two player games of perfect information. This include any game where each player moves at most once on any path.

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

Let X denote the set of non-terminal nodes, that are partioned into the X_1 and X_2, the decision nodes of the two players. Any pure behavior strategy profile σ induces a terminal node starting at any non-terminal node x. Write $u(\sigma(x))$ for the payoffs at the terminal node so induced. For any $x \in X$, $\bar{\sigma}(x)$ denotes the unique backwards induction path induced by $\bar{\sigma}$ starting at x, and $u(\bar{\sigma}(x))$ denotes the payoff allocation at the corresponding terminal node. Given any terminal node $z \in Z$, let $\phi(z)$ denote the path from the initial node x_0 to z.

Definition 7 *Fix a terminal node z, a path $\phi(z)$, and a node x on this path, where player i moves. Player i has an incentive to deviate at x if $u_i(\bar{\sigma}(x)) > u_i(z)$. Player i has an incentive to deviate from $\phi(z)$ if there exists a node x on this path where he has an incentive to deviate.*

Remark 8 *No player has an incentive to deviate from the path to $\bar{z}$, the backwards induction outcome. If $z \neq \bar{z}$, then some player has an incentive to deviate from the path to z.*

We restrict attention in this paper to the sustainability of outcomes that Pareto dominate the backward induction outcome $\bar{z}$, in a class of two-player games. More precisely, consider pairs (Γ, z^*), where Γ is a generic two-player game and z^* is a terminal node, where $u_1(z^*) > u_1(\bar{z})$ and $u_2(z^*) > u_2(\bar{z})$. We assume that the pair (Γ, z^*) satisfy the following assumption:

Assumption 9 *Only one player has an incentive to deviate on the path to z^*.*

Let d index the player who has and incentive to deviate, and let j index the player who does not have an incentive to deviate.

We now show that the class we consider includes every outcome that Pareto-dominates the backwards induction outcome in *games where each player moves at most once.* We define these to be generic two-player games of perfect information where along any path, each player moves at most once. Without loss of generality, player 1 moves at the initial node, choosing from a finite set A_1 and player 2 moves after some choices of player 1.

Lemma 10 *If Γ is a game where each player moves at most once, and z^* Pareto dominates $\bar{z}$, then only player 2 has an incentive to deviate on the path to z^*. Furthermore, $u_1(\bar{\sigma}(a_1^*)) < 0$.*

Proof. Denote the path to z^* from the initial node by a_1^*, a_2^*. We claim that at a_1^*, player 2 has *an incentive to deviate*, i.e. his optimal action differs from a_2^*; if this was not the case, then since player 1 prefers z^* to $\bar{z}$, z^* would be the unique backwards induction outcome. To see that player 1 does not have an incentive to deviate at the initial node x_0, note that $u_1(\bar{\sigma}(x_0)) = 0 < u_1(z^*)$. To prove the second part, if $u_1(\bar{\sigma}(a_1^*)) > 0$, then 1 would choose a_1^* under the backwards induction strategy profile, contradicting our assumption that $\bar{z}$ was the backwards induction outcome. ■

Let the pair (Γ, z^*) satisfy our assumption, that only one player, indexed by d, has the incentive to deviate. Let $X_i(z^*) = \{x \in X_i : x < z^*\}$ denote the set of decision nodes of player i on the path $\phi(z^*)$ and let Let $X_i^C = X_i - X_i(z^*)$ denote the complement. Let $\hat{X}_d(z^*)$ denote the set of decision nodes of player d where he has an incentive to deviate from $\phi(z^*)$, i.e. the set $\{x \in X_d(z^*) : (\bar{\sigma}(x))) > u_d(z^*)\}$.

Let $\hat{x}_d$ denote the node where d's incentive to deviate is maximal, i.e.

$$\hat{x}_d = argmax_{x \in X_i(z^*)}(u_d(\bar{\sigma}(x))).$$

Let $u_d(\bar{\sigma}(x))) = 1 + g$, where $g > 0$. Recall that $u_d(z^*)$ was normalized to 1.

The following assumption ensures that the beliefs at unreached information sets in the equilibrium we construct satisfy the D1 criterion.

Assumption 11 *$\hat{x}_d$ precedes any other node in $\hat{X}_d(z^*)$, i.e. if $x \in \hat{X}_d(z^*), \hat{x}_d \leq x$.*

The information structure in the repeated game is as follows. First, partition the set of outcomes in the stage game so that D denotes a terminal node that arises after a deviation by d from $\phi(z^*)$ at a node where she has an incentive to deviate. Let N denote the complement, $N = Z - D$. Player d's observed history is denoted B if there is any instance of D in any of the last K periods; otherwise, her observed history is denoted G. Player j observes whether d's history is G or B before the players play the stage game Γ. Player d observes no information regarding the past play of player j.

In order to define repeated game strategies, we first define the following strategies/strategy profiles in the game Γ.

Let $\sigma^* = (\sigma_d^*, \sigma_j^*)$ denote the strategy in Γ where $\phi(z^*)$ is played unless some player deviates from $\phi(z^*)$, in which case players continue with $\bar{\sigma}$.

Define the strategy $\hat{\sigma}_i$ in Γ as follows.

$$\hat{\sigma}_i(x) = \begin{cases} \sigma_i^*(x) & \text{if } x \in X_i(z^*) \text{ and } x \geq \hat{x}_d \\ \bar{\sigma}_i(x) & \text{otherwise.} \end{cases}$$

The repeated game strategies are as follows.

The players play σ^* at G.

Player j plays $\hat{\sigma}_j$ at B.

Player d plays $\bar{\sigma}_d$ at B if $m > m^*$ and plays $\hat{\sigma}_d$ at B if $m < m^*$.

Define $\tilde{x}_j$ as the maximal element under $(\leq)$, precedence, of the set $\tilde{X}_j$, where

$$\tilde{x}_j = \{x \in X_j(z^*), x \leq \hat{x}_d, u_j(\bar{\sigma}(x)) > u_j(\bar{\sigma}(\hat{x}_d)\}.$$

We show that the $\tilde{X}_j$ is non-empty, so that $\tilde{x}_j$ is well defined. If $\tilde{X}_j$ is empty, this implies that for all $x < \hat{x}_d, x \in X_j(z^*), u_j(\bar{\sigma}(x)) \leq u_j(\bar{\sigma}(\hat{x}_d))$. Since player d's incentive to deviate from $\phi(z^*)$ is maximal at $\hat{x}_d$, $u_d(\bar{\sigma}(x)) \leq u_d(\bar{\sigma}(\hat{x}_d))$ if $x < x, x \in X_d$. These two facts imply that $\bar{\sigma}(\hat{x}_d)$ is the backward induction outcome, $\bar{z}$. But since $u_d(\bar{\sigma}(\hat{x}_d) > u_d(z^*)$, this contradicts the assumption that z^* Pareto-dominates $\bar{z}$. We may therefore define

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

Notice that ℓ is player j's loss from continuing on the path $\phi(z^*)$ at $\tilde{x}_j$ if player d continues with his backward induction strategy, and we have established that $\ell > 0$.

We now verify the optimality of these repeated game strategies.

Consider signal G. If $K > \bar{K}$, then player d has no incentive to deviate from σ^* at G. Given this, neither does player j, since by assumption, j does not have an incentive to deviate from $\phi(z^*)$.

Now consider signal B. Consider any node $x < \tilde{x}_j$ on $\phi(z^*)$. Given that player j plays $\bar{\sigma}_j$ at $\tilde{x}_j$, backward induction establishes that $\bar{\sigma}_i$ is optimal for $i \in \{1, 2\}$ at such node x that precedes $\hat{x}_d$.

Consider next the node $\hat{x}$. If player d plays $\bar{\sigma}_d$ at this node, the specified continuation strategies imply that she gets a current payoff of $u_d(\bar{\sigma}(\hat{x}_d))$ and a continuation value of $V^K(K)$. If instead she continues on path $\phi(z^*)$, she gets a payoff of $u_d(z^*)$ and a continuation value of $V^K(m-1)$. Recall that g was defined by

$$g := u_d(\bar{\sigma}(\hat{x}_d)) - u_d(z^*).$$

Thus the payoff difference between these two choices equals

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

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

There remains the critical node, $\tilde{x}_j$. Let us now consider player j's payoff difference between continuing on the path $\phi(z^*)$ at $\tilde{x}_j$, and playing the backward induction strategy. Let θ denote the probability assigned by j to the type of player d being $m > m^*$. The payoff difference between continuing on $\phi(z^*)$ and playing the backward induction strategy equals

$$(1-\theta)u_j(z^*) + \theta u_j(\bar{\sigma}(\hat{x}_d)) - u_j(\bar{\sigma}(\tilde{x}_j)).$$

Since player j has no incentive to deviate, $u_j(z^*) > u_j(\bar{\sigma}(\tilde{x}_j))$. Also, $\ell = u_j(\bar{\sigma}(\tilde{x}_j)) - u_j(\bar{\sigma}(\hat{x}_d) > 0$. Thus this payoff difference is negative as long as

$$\theta > \frac{u_j(z^*) - u_j(\bar{\sigma}(\hat{x}_d))}{u_j(z^*) - u_j(\bar{\sigma}(\tilde{x}_j)} := \bar{\theta} \in (0,1). \tag{19}$$

Lemma 12 *Suppose that* $\exists x \in X_d^*, x < \tilde{x}_j : u_d(\bar{\sigma}(x)) > u_d(\bar{\sigma}(\tilde{x}_j))$*. Then D1 implies* $\theta(\tilde{x}_j) = 1$.

Thus, if the game is such that player d moves on the path $\phi(z^*)$ before $\tilde{x}_j$, and his backward induction strategy does not specify continuing along the path, then at $\tilde{x}_j$, player j assigns probability one to player d deviating from $\phi(z^*)$ at $\hat{x}_d$, and inequality 19 is satisfied. In this case, we have an equilibrium as long as $K \geq \bar{K}$. If the game does not satisfy this condition, then we need $\frac{\lfloor m^* \rfloor}{K}$ to be small enough so that it is less than $1 - \bar{\theta}$. This is always ensured if K is large enough.

References

Banks, J. S., and J. Sobel (1987): "Equilibrium Selection in Signalling Games," *Econometrica*, 55, 647–661.

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.

Cripps, M. W., G. J. Mailath, and L. Samuelson (2004): "Imperfect Monitoring and Impermanent Reputations," *Econometrica*, 72(2), 407–432.

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.

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.

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 (2015): "Observations on cooperation," *Available at SSRN 2558570*.

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

Kandori, M. (1991): "Cooperation in Finitely Repeated Games with Imperfect Private Information," Princeton University.

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

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

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

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

Padilla, A. J., and M. Pagano (1997): "Endogenous communication among lenders and entrepreneurial incentives," *Review of Financial Studies*, 10(1), 205–236.

Pagano, M., and T. Jappelli (1993): "Information sharing in credit markets," *The Journal of Finance*, 48(5), 1693–1718.

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

Sugaya, T., and A. Wolitzky (2014): "Perfect Versus Imperfect Monitoring in Repeated Games," Discussion paper, mimeo.

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

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