The Design of Credit Information Systems*

V Bhaskar† Caroline Thomas‡

February 8, 2017

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. Borrowers who default should be temporarily excluded in order to efficiently incentivize repayment. However, lenders have an incentive to lend to borrowers who are near the end of their exclusion period, undermining this punishment. With perfect information on past behavior, no lending can be sustained, independent of the length of memory. Coarse information about borrowers' histories improves outcomes, by mitigating lender moral hazard. By pooling recent defaulters with those who defaulted further in the past, coarse information induces borrower adverse selection. This disciplines lenders, who are now unable to target loans towards those defaulters who are least likely 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 can also improve efficiency, by raising the proportion of likely re-offenders in the pool of recent defaulters. To summarize, endogenously generated borrower adverse selection mitigates lender moral hazard.

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

Keywords: credit markets, one-sided moral hazard, information design, repeated games with randomly matched players.

---

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

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

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

1 Introduction

We examine the design of information and rating systems in large markets where transactions are bilateral and moral hazard is one-sided. Many examples fit: 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. Our leading example is unsecured debt — the borrower takes a loan and must subsequently decide whether to repay or willfully default. The market is large and each pair of agents transacts infrequently. Thus, opportunistic behavior (by the seller, builder or borrower) can be deterred only by a "reputational mechanism", whereby opportunism results in future exclusion. We assume that information on past transgressions is subject to *bounded social memory* and is retained only for a finite length of time. While this is plausible in any context, it is legally mandated in many credit markets. In the United States an individual's bankruptcy record 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. Bounded memory also arises under policies used by internet platforms to compute the scores summarizing their participants' reputations. For example, Amazon lists a summary statistic of seller performance over the past 12 months — given that buyers have limited attention, this may serve as to effectively limit memory. In the United States, 24 states and many municipalities[1] have introduced "ban the box" legislation, prohibiting employers from asking job applicants about prior convictions unless those relate directly to the job.[2]

How do markets function when moral hazard is important and information systems are constrained by bounded memory? For concreteness, consider the credit market example.[3] Since borrowers have a variety of sources of finance in a modern economy, we study a model with a large number of long-lived borrowers and lenders, where each borrower-lender pair interacts only once. 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,

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

[2] One should note 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 may compel online search engines to delete past records pertaining to them.

[3] We abstract from many realistic features of such markets in order to focus on a key problem.

and has short-term incentives to willfully default. Additionally, there is a small chance of involuntary default.[4] Thus lending can only be supported via long-term repayment incentives whereby default results in the borrower's future exclusion from credit markets.

In our large-population random matching environment, however, 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 past transgressions. Thus a borrower can only be deterred from willful default if this record indicates that she is likely to default on a subsequent loan.[5] With bounded memory, *disciplining lenders* to not lend to borrowers who have defaulted recently turns out to be a non-trivial problem. What are the information structures and strategies that support efficient lending in such an environment?

A natural conjecture is that providing maximal information is best, so that the lender has complete information on the past K outcomes of the borrower, where K is the bound on memory. This turns out to be false. Perfect information on the recent past behavior of the borrower, in conjunction with bounded memory, precludes any lending, as it allows lenders to cherry-pick those borrowers with the strongest long-term incentives to repay. In *any* equilibrium that satisfies a mild requirement of being robust to small payoff shocks, borrower exclusion unravels.[6]

In particular, a borrower whose most recent default is on the verge of disappearing from her record has the same incentives as a borrower with a clean record. Thus she will repay a loan if long-term incentives are such that a borrower with a clean record does so. Lenders, who can distinguish her from more recent defaulters, find it profitable to extend her a loan, 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 problem is that, under perfect information, lenders cannot be disciplined to not make loans to borrowers with a bad record.[7] In our context, rogue lending undermines borrowers' long-term incentives, as

[4] Willful, or "strategic" default is particularly attractive in the case of no-recourse loans. (If the borrower defaults, the lender can seize the collateral but the borrower is not liable for any further compensation in case the value of the collateral does not cover the full value of the defaulted amount.) A large fraction of commercial mortgages are non-recourse. These are frequently re-negotiated if the value of the asset plummets, the threat of strategic default giving the lender some bargaining-power vis-à-vis the lender.

[5] The reader may ask why the lender cannot be disciplined by allowing future borrowers to condition their behavior on the lender's current decision. This mechanism, which is standard in many repeated games, turns out to be unviable given the informational constraints of our context. This is discussed in see Section 3.3.

[6] The payoff shocks imply that agents have strict incentives at every information set, except possibly for a set of shock realizations that is negligible. They ensure that the equilibria we consider are robust.

[7] Reckless sub-prime mortgage lending has been at the heart of the recent financial crisis. The expectation that house prices and incomes would rise, together with the emergence of collateralized debt obligations inflating the demand for mortgage-based financial products, meant even loans to borrowers with low ("sub-

lenders seeking profit opportunities impose negative externalities on other lenders.

The negative result when lenders have perfect information about bounded borrower histories leads us to explore information structures that provide the lender with coarse information about these histories.[8] Specifically, the lender is told only whether the borrower has ever defaulted in the past K periods (labelled a *bad credit history*) or not (labelled a *good credit history*). Although all agents have the same short-term incentive to default and are identical in their probability of involuntary default, their long-term incentives to repay a current loan now differ according to the most recent instance of default in their 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 adverse selection among the pool of borrowers with a bad credit history, thereby mitigating lender moral hazard.[9] Our question is, how can a coarse information structure be tailored to sustain efficient outcomes?

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

Efficiency can be improved, somewhat paradoxically, by a non-monotonic information partition, where borrowers with multiple defaults are treated favorably and pooled with non-defaulters. This provides strong incentives for defaulters to re-offend. However, such a mechanism may be vulnerable to bribery, and we focus our attention on monotone information structures that are immune to gaming. These information structures turn out to be the simple binary information structures described above, with borrowers being partitioned into good and bad credit risks.

We show that even if loans are very profitable, efficiency can be achieved in an equilibrium where borrowers with bad credit histories are provided loans with positive probability.

prime") FICO scores were considered worthwhile by lenders.

[8]Coarse information is widely used in many markets. Amazon provides summary statistics on sellers via a five-star rating system. The finer FICO scores are often bundled into "sub-prime" (620 FICO and below), "near prime" (621 - 679), and "prime" (680 or greater) ratings. (Silvia (2008))

[9]In Section 3.1.2 we show that our results remain valid in the presence moderate exogenous borrower adverse selection, as this makes lending to a borrower with a bad credit history less attractive for the lender.

Some of them will default, altering the constitution of the pool of borrowers with bad credit histories, as borrowers with stronger incentives to re-offend will be over-represented. This serves to discipline lenders. Paradoxically, if individual loans are very profitable, an equilibrium with random exclusion may result in low overall profits for lenders, by inducing a large pool of borrowers with bad records.

Finally, we show that mandates preventing lenders from "chasing borrowers", namely by requiring an initial loan application by the borrower, can increase efficiency, and indeed increase lender profits. In conjunction with coarse information, such a rule transforms the interaction between borrowers and lenders into a signaling game. Among the borrowers with a bad credit history, those who intend to default have stronger incentives to apply than those who intend to repay. This causes lenders to be suspicious of applicants with a bad record, and dissuades them from lending.

To summarize the main insight of our paper: Although moral hazard is one-sided — on the part of borrowers in our leading example — the key problem is in ensuring that lenders do not lend to recent defaulters. Providing lenders with perfect information about borrowers' recent past behavior leads to rogue lending, undermining the credit market altogether. A binary coarse information structure endogenously produces borrower adverse selection, and can be used to discipline the lender. We show that efficient lending can be supported under a simple, binary information structure.

The remainder of this section discusses the related literature. Section 2 sets out the model and derives efficiency benchmarks, which can be attained with infinite memory. It also shows that with bounded perfect memory, no lending can be supported. Sections 3 and 4 contain the main results of this paper. They show that a simple, coarse information structure can support efficient lending. In section 5 we show that preventing lenders from chasing borrowers improves lender incentives and excessive lending. In section 6 we show that our main results apply to a large class of two-player (stage) games. Section 7 concludes.

1.1 Related Literature

There is a growing literature on the role of information in the functioning of credit market. Musto (2004) empirically documents the effects of expunging defaults, and argues that this legislation has adverse consequences for resource allocation, by removing information from the market. Chatterjee, Corbae, Nakajima, and Ríos-Rull (2007) set out a quantitative model of default, where default records are stochastically expunged, that replicates the main empirical features of the market for unsecured debt in the US. Notably, their model does not incorporate adverse selection, suggesting that residual private information is small in modern credit markets.

Contrary to Musto (2004), there is also a literature that argues that limited records may be welfare improving in the presence of adverse selection. Elul and Gottardi (2015) analyze a market with two types of borrower, who are also subject to moral hazard. Utilitarian efficiency dictates lending to both types, but lending to the high risk borrower is unprofitable. Bounded memory can increase efficiency, by allowing high risk borrowers to pool with low risk ones. Kovbasyuk and Spagnolo (2016) consider 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. Padilla and Pagano (1997) also argue that information provision may be excessive. Our paper differs from this literature on two dimensions. First, we focus on the moral hazard, rather than adverse selection — given the information available to lenders on consumer characteristics, it could be argued the residual private information is small. Second, we model the expunging of records as deterministic rather than stochastic. This is consistent with the legal requirement.

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), Deb (2008) and many others. Since we assume a large (continuum) population, this precludes the use of contagion strategies that are usually used in this literature. 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 stage-game — almost exclusively the prisoner's dilemma[10] — 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, our stage game is generic in the extensive form, but its strategic form counterpart would be non-generic.

An alternative application of our model is the interaction between a buyer and a seller, where the buyer makes a purchase decision, and the seller must decide what quality to supply. This problem has been studied in the large literature on seller reputation. Most closely related is Liu and Skrzypacz (2014), which assumes that buyers are short-lived and have bounded information on the seller's past decisions, but do not observe the magnitude of past sales. They demonstrate that equilibrium displays a cyclical pattern, whereby the

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

seller builds up his reputation before milking it.[11] 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 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, and also shows that our equilibria are robust. 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 sustain greater efficiency than memory can.

Our work also relates to the recent interest in information design, initiated by Kamenica and Gentzkow (2011), and pursued by Bergemann and Morris (2016), Taneva (2016), Mathevet, Perego, and Taneva (2016) and Ely (2017). While this literature has focused on the case of one or few players, our design question relates to a large society. Whereas the distribution of types or states is usually exogenous in the information design context, the induced distributions over types (or histories) in the present paper arise endogenously as a result of 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 positive analysis.[12] In each period, individuals from population 1 (the lenders) are

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

[12] As we will see in Section 3.3, incentives for the lender have to be provided within the period.

randomly matched with individuals from population 2 (the borrowers),[13] to play the following sequential-move game, illustrated in Figure 1a. First, the lender (he) chooses between $\{Y, N\}$, i.e. whether or not to extend a loan. If he chooses N, the game end, with payoffs $(0, 0)$. If he chooses Y, then the borrower (she) invests this loan in a project with uncertain returns. With a small probability λ, the borrower is unable to repay the loan, i.e. she is constrained to default, D. With the complementary probability she is able to repay the loan, and must choose whether or not do so, i.e. she 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 occurs. When she is unable to repay so that her default in involuntary, her payoff is 0. If she willfully defaults even though she is able to repay, her payoff is $\pi_b + \tilde{g}$, where $\tilde{g} > 0$. Suppose the borrower chooses R when she is able to. 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 Figure 1a. 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 willfully defaults are $(-\ell, 1 + g)$.[14] The associated strategic form of the stage-game, given in Figure 1b, is a one-sided prisoner's dilemma. We assume that only the borrower can observe whether or not she is able to repay, i.e. the lender or any outside observer can only observe the outcomes in the set $O = \{N, R, D\}$.[15]

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

[13]If a borrowing opportunity only arises in each period with some probability, this affects the borrower's effective discount factor. This makes no qualitative difference, except if defaults and not getting a loan are not distinguishable in the long run; see Section 3.6.

[14]For the loans considered in this paper, short-term defaulting costs considered in the consumer finance literature (e.g. filing costs, arbitration costs, social stigma — see Fay, Hurst, and White (2002)) are not sufficient to dissuade borrowers from defaulting. Loans no borrower would ever default on are implicitly subsumed in the N outcome. Thus, we study the sustainability of loans, the repayment of which requires long-term incentives.

[15]An alternative specification of the model is as follows. The borrower has two choices of project. The safe project results in a medium return M 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, $1 - \theta < 1 - \lambda$, and a zero return with complementary probability. Let $(1 - \lambda)M - r = 1$,$(1 - \theta)H - r = 1 + g$, $(\theta - \lambda)r = (1 + \ell)$. Then 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). Our analysis can be generalized to this case, but we do not pursue this here.

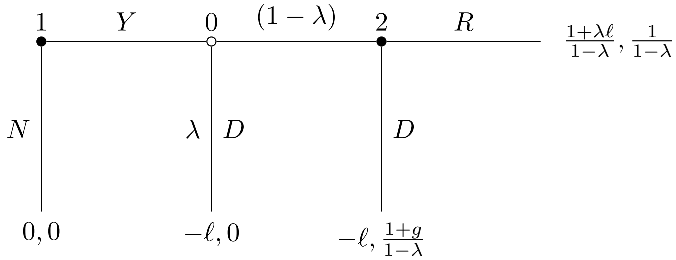

(a) Extensive form.

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

(b) Strategic form.

Figure 1: Extensive and strategic form representations of the stage-game

in every period unless future lenders have information about her behavior. We henceforth assume that they do — the details of the information structure will vary, depending on the context we consider.

We focus on stationary Perfect Bayesian Equilibria, where agents are sequentially rational at each information set, with beliefs given by Bayes rule wherever possible. The stationarity assumption implies that players do not condition on calendar time. Furthermore, we shall focus on equilibria where all lenders follow the same strategy, and all borrowers follow the same strategy. Finally, we will also require that our equilibria are purifiable, as we now explain.

2.1 Payoff Shocks: The Perturbed Game

We now describe 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 at $x \in X$, 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 $a_k \in A(x)$, that player's payoff from action a_k is augmented by εz_x^k, where $\varepsilon > 0$. The scalar z_x^k is the k^{th} component of z_x, where $z_x \in \mathbb{R}^{|A(x)-1|}$ is the realization of a random variable with bounded support. We assume that the random variables $\{Z_x\}_{x \in X}$ are independently distributed, and that their distributions are atomless. Player $\iota(x)$ observes the realization z_x of the shock before being called upon to move. In the repeated version of the perturbed game, $\Gamma^\infty(\varepsilon)$, we assume that the shocks for any player are independently distributed across periods.[16] 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 Γ^∞ 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 *sequentially 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. The following lemma, proved in appendix A.6, is useful in our analysis:

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

2.2 The Infinite Memory Benchmark

Suppose that each lender can observe the entire history of transactions of the borrower. That is, a lender matched with a borrower at date t observes the outcome of the borrower in periods $1, 2, .., t-1$. Assume that the borrower is patient enough, so that there exists an equilibrium where lending takes place.[17] We now set out some efficiency benchmarks, while postponing the precise details of the equilibrium strategies to Section 3.5.

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 $(1-\delta)g$.[18] 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, willful or involuntary, she is excluded from the lending market for K periods. The incentive constraint ensuring that she prefers repaying when able is then

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

where $V^K(0)$ denotes her payoff when she is in good standing, and $V^K(K)$ her payoff at the

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

[17] The precise condition is $g < \frac{\delta(1-\lambda)}{1-\delta(1-\lambda)}$.

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

beginning of the K periods of punishment. These are given by

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

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

The most efficient equilibrium in this class has K large enough to provide the borrower incentives to repay when she is in good standing, but no larger. Call this value $\bar{K}$, and assume that the incentive constraint (1) holds as a strict inequality when $K = \bar{K}$ — this assumption will be made throughout the paper, and is satisfied for generic values of the parameters (δ, g, λ). The payoff of the borrower when she is in good standing is $\bar{V} := V^{\bar{K}}(0)$, i.e. it is given by equation (2) with $K = \bar{K}$.

A more efficient equilibrium can be sustained if players observe the realization of a public randomization device at the beginning of each period, and this is a part of the public history. The best equilibrium payoff for the borrower is one where the incentive constraint just binds, and equals

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

Notice that in both types of equilibria, the lender obtains a payoff of 1 when she meets a borrower in good standing, and a payoff of 0 on meeting a borrower who should be excluded. Thus the expected payoff for the lender increases as punishments are reduced, and, to this extent, there is a commonality of interests across equilibria.[19]

2.3 Perfect Bounded Memory

Henceforth, we shall assume that lenders have bounded memory, i.e. we assume that at every stage, the lender observes a bounded history of length K of past play of the borrower she is matched with in that stage. 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.

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.

[19] The arguments in Fudenberg, Kreps, and Maskin (1990) can be used to support a borrower payoff higher than V^*, if a borrower in good standing has access to a randomization device that allows her to willfully default with some probability. However, this argument requires that the realizations of the randomization device also be a part of the infinite public record, imposing too great an informational burden for our purposes.

Proposition 2 *Suppose that $K \geq 2$ is arbitrary and the lender 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 who receives a loan in period t knows that the lender at $t + 1$ cannot observe the outcome at $t - 1$. The payoff shocks in $\Gamma(\epsilon)$ imply that for any strategy of the lender at $t + 1$, at t the borrower is indifferent between R and D only on a set of measure zero, and has a strict best response for almost all realizations of the payoff shock. Therefore, the borrower cannot condition her behavior at t 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 is therefore also unable to condition his behavior at t on the outcome at date $t - 1$. Thus, neither player can condition their behavior on the past history. The argument for the case when $K \geq 2$ 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.

The following aphorism encapsulates the message of this section: *What is eventually forgotten never happened.* As we shall see in section 3.6, this message also applies to alternative formulations of forgetting defaults, and also implies that one cannot provide incentives by conditioning on the behavior of the borrower (see section3.3).

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

3 Coarse Information on Borrowers

We now consider information structures where the lender has partial information on the outcomes of the lender. A deterministic information structure consists of a finite signal space S and a mapping $\tau : O^K \to S$. More simply, it consists of a partition of the set of K-period histories, O^K, with each element of the partition being associated with a distinct signal in S. Therefore, we will call this a partitional information structure. Contrast it with a random information (or signal) structure, where the range of the mapping is the set of probability distributions over signals, so that $\tau : O^K \to \Delta(S)$. Note that in both cases, the signal *does not* depend upon past signal realizations, since otherwise one could smuggle in infinite memory on outcomes.

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

We 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 equilibria with temporary exclusion.

3.1 A Simple Information Structure

Let us assume that the length of memory is $K \geq \max\{\bar{K}, 2\}$. Suppose that the information structure is the following binary partition of O^K. The lender observes a "bad credit history" signal B if and only if the borrower has had an outcome of D in the last K periods, and observes a "good credit history" signal G otherwise. This information structure is *monotone* in the sense that if bounded histories are ordered by default rates then the information structure does not pool worse histories with better ones while separating intermediate ones. Since this information structure will recur through this paper, it will be convenient to label it *the simple information partition/structure.*

The borrower has complete knowledge of her own private history, since she knows the entire history of past transactions. Information on events that occurred more than K periods ago is irrelevant, since no lender can condition on it. Under the simple information structure, the following partition of K-period private histories will be used to describe the borrower's incentives. Partition the set of private histories, into $K+1$ equivalence classes, indexed by m. More precisely, let t' denote the date of the most recent incidence of D in the borrower's history, and let $j = t - t'$, where t denotes the current period. Define $m := \min\{K+1-j, 0\}$. Under the simple information structure, if $m = 0$ the lender observes G while if $m \geq 1$ the lender observes B. Thus, m represents the number of periods that must elapse without default before the borrower gets a good history. When $m \geq 1$, this value is the borrower's private information. In particular, among lenders with a bad credit history, the lender is not able to distinguish those with a lower m from those with a higher m.

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

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

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

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

In Appendix A.1 we show that (6) is satisfied whenever $K \geq \bar{K}$.

Now consider the incentives to repay for a borrower with an arbitrary $m > 1$. Choosing D over R entails the payoff difference

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

We have seen that when $m = K$, the above expression is positive, while when $m = 1$, the expression is negative. Thus there exists a real number, denoted $m^*(K) \in (1, K)$, that sets the payoff difference equal to zero. We assume that $m^*(K)$ is not an integer, as will be the case for generic payoffs.[20] If $m > m^*(K)$, the borrower chooses D when offered a loan. If $m < m^*(K)$, she chooses R. Intuitively, a borrower who is close to getting a clean history will not default, just as a convict nearing the end of his sentence has incentives to behave.

We now describe the beliefs of the lender when he observes B. In every period, the probability of involuntary default is constant, and equals λ. Furthermore, under the candidate strategy profile, a borrower with a bad credit history never gets a loan and hence transits deterministically through the states $m = K, K-1, .., 1$. Therefore, the invariant (steady state) distribution over values of m induced by this strategy profile gives equal probability to each of these states.[21] Consequently, the lender's beliefs regarding the borrower's value of m, conditional on signal B, are given by the uniform distribution on the set $\{1, \ldots\ldots, K\}$. He therefore attributes probability $\lfloor m^*(K) \rfloor / K$ to a borrower with signal B repaying a loan, where $\lfloor x \rfloor$ denotes the largest integer no greater than x. Simple algebra shows that making a loan to a borrower with a bad credit history is strictly unprofitable for the lender if

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

Suppose that ℓ is large enough that the lender's incentive constraint (8) is satisfied. Then he finds it strictly optimal not to lend after B, and to lend after G. We have seen that, since $K \geq \bar{K}$, it is optimal for the borrower to repay when she has a clean history G. Moreover, if granted a loan after B, she has strict incentives to repay as long as $m < m^*(K)$ and to default if $m > m^*(K)$. Thus, there exists an equilibrium that is sequentially strict, and is therefore purifiable. In other words, providing the borrower with coarse information, so that he does not observe the exact timing of the most recent default, overcomes the impossibility result in Proposition 2. Even though those types of borrowers who are close to "getting out of jail" would choose to repay a loan, the lender is unable to distinguish them from those whose sentence is far from complete. He therefore cannot target loans to the former. In other words, coarse information endogenously generates adverse selection that mitigates *lender moral hazard*.

It remains to investigate the conditions on the parameters that ensure that the incentive constraint (8) is satisfied. This depends on $m^*(K)$, whose properties are summarized in the

[20]This ensures that the equilibrium is sequentially strict, permitting a simple proof of purifiability.
[21]The invariant distribution $(\mu_m)_{m=0}^K$ has $\mu_0 = \frac{1}{1+K\lambda}$ and $\mu_1 = \cdots = \mu_K = \frac{\lambda}{1+K\lambda}$.

following lemma.

Lemma 3 *If $K = \bar{K}$, then $m^*(K) \in (1, 2)$. Moreover, $\frac{\lfloor m^*(K) \rfloor}{K} \to 0$ as $K \to \infty$.*

The second part of the above lemma, in conjunction with our previous discussion, immediately implies:

Proposition 4 *An equilibrium where the lender lends after observing G and does not lend after observing B exists as long as K is sufficiently large. Such an equilibrium is sequentially strict and therefore purifiable.*

Note however that if $K > \bar{K}$, then the equilibrium results in inefficiently long borrower exclusion relative to the length required for overcoming borrower moral hazard.

3.1.1 Efficiency When Loans Are Not Too Profitable

We now investigate the conditions under which an equilibrium with punishments of minimal length, $\bar{K}$, exists, assuming that the memory length also equals $\bar{K}$, for $\bar{K} \geq 2$. Lemma 3 shows that $m^*(\bar{K}) \in (1, 2)$, so that a borrower with type $m = 1$ repays, while all other types $m > 1$ default, giving rise to a steady state repayment probability of $\frac{1}{\bar{K}}$. The lender has strict incentives not to lend to a borrower with a bad credit history if

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

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

So far, we have restricted attention to partitional information structures on O^K. We now show that random signal structures can increase 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

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

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

so that

$$V^{\bar{K},x}(0) = \frac{1-\delta}{1-\delta(1-\lambda+\lambda\,\delta^{\bar{K}-1}(x\delta+1-x))}. \tag{10}$$

The borrower agrees to repay at G if and only if

$$(1-\delta)g \leq \delta(1-\lambda)\left(1-\delta^{\bar{K}-1}(x\delta+1-x)\right)\ V^{\bar{K},x}(0). \tag{11}$$

We show in Appendix A.2.1 that the right hand side of (11) is a strictly increasing function of x. Moreover, by the definition of $\bar{K}$, (11) is satisfied when $x=1$, and violated when $x=0$. There therefore exists $x^*(\bar{K}) \in (0,1]$ such that (11) is satisfied for every $x \in [x^*(\bar{K}),1]$.[22]

It remains to examine whether this modification preserves the incentive of lenders not to lend at B. In Appendix A.2.2 we show that, 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 random information structure will do better than the simple information structure. In particular, whenever there is a $\bar{K}$ memory partitional information structure that supports an equilibrium where the borrower's payoff is $\bar{V}$, a non-partitional signal structure such as the random one set out here can approximate V^*, the efficient borrower payoff defined in (4).

Proposition 5 *Suppose $\bar{K} \geq 2$. If loans are not too profitable, so that $\frac{1}{\bar{K}} < \frac{\ell}{1+\ell}$, there exists a monotone information partition and an equilibrium that achieves the borrower payoff $\bar{V}$. In this case, there also exists a random signal structure and an equilibrium that achieves the efficient payoff V^*.*

Remark 6 *The lender's incentive compatibility constraint under a random signal structure can be weakened. The precise condition is: $\frac{x^*}{x^*+\sum_{m=2}^{\bar{K}}(1-\lambda)^{m-1}} < \frac{\ell}{1+\ell}$.*

One concern with non-partitional information structures, such as the random structure discussed in this section, is that they may be vulnerable to manipulation. Since the report

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

in the final period of the punishment phase is random, influential borrowers may be able to persuade the credit agency (information provider) to modify the report in their favor. This would still be consistent with the provider having a fraction x of reports at $m = 1$ being B. When manipulation is possible, an influential borrower's incentive constraint would be violated, since the effective length of her punishment phase would only be $\bar{K} - 1$.

3.1.2 Adverse Selection

In focusing on borrower moral hazard, we have assumed away borrower adverse selection. We now show that moderate adverse selection among borrowers may actually improve matters, by mitigating lender moral hazard. In other words, exogenous borrower adverse selection can augment the adverse selection that is endogenously generated by the coarse information on borrower histories.

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

3.1.3 Varying Loan Terms

We have assumed that the repayment terms of loans are the same, regardless of the credit history of the borrower. How would our analysis be altered if we allowed the lender to charge a higher interest rate from a borrower with a bad record? A higher interest rate would make such a loan more profitable, conditional on repayment; however, it also increases the borrower's incentive to default. Nevertheless, a revealed preference argument implies that the individual lender must be better off from being able to tailor interest rates to the borrower's record. Thus, the first order effect is to undermine borrower exclusion.

One may conjecture that since a defaulting borrower is subject to higher interest rates, this serves as an alternative form of punishment. However, this cannot be the case: higher interest rates alone cannot serve as a sufficient punishment, since a defaulter who is charged higher rates can default once again. Thus borrower exclusion is necessary, and allowing for variable terms does not significantly affect the analysis.

3.2 Non-monotone Information Structures

Recall that $\bar{K}$ periods of exclusion are sufficient to provide incentives for the borrower not to willfully default. Nonetheless, the equilibria constructed in the previous section may well need longer periods of exclusion. In particular, if loans are sufficiently profitable because ℓ is small, then $\bar{K}$ periods of punishment may not suffice to discipline the lender, even though they suffice to discipline the borrower.

How can we discipline lenders, given that their incentives cannot depend upon future play, but only upon borrower behavior in the current period? One idea is as follows: suppose that we reward borrowers for defaulting on a lender who should not have made a loan. This makes it more likely that a borrower with a bad credit history will default, and makes lending to them less attractive. We can implement this idea, under bounded memory, by assigning a good credit history to borrowers with *two* defaults, while those with a single default are assigned a bad credit history. We call such an information structure non-monotone, since borrowers with no defaults are pooled with borrowers with two defaults, while one-default borrowers are excluded from this pool.

There still remains the problem identified in Section 2, that the incentives of a borrower with a single default which occurred exactly K periods ago (i.e. a borrower with $m = 1$) are identical to that of a borrower with no defaults. Thus in any purifiable equilibrium, the behavior of the two borrowers must be identical, and if the borrower with no defaults repays for sure, so must the borrower with a single default and $m = 1$. This suggests that the probability of repayment of a lender with a bad credit history cannot be reduced below $\frac{1}{K}$. Nonetheless, the following proposition constructs an information structure where this repayment probability is zero. The trick here is to allow for memory that is one period longer than is required for borrower incentives, so that $K = \bar{K} + 1$. The additional period is not used to punish the borrower, but instead to retain information for disciplining the lender.

Proposition 7 *For any $\ell > 0$, there exists a non-monotone informational structure, and an associated sequentially strict equilibrium that can approximate $\bar{V}$.*

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

Remark 8 *The proposition applies also when $\bar{K} = 1$ — in contrast with Proposition 5, which does not.*

Can the efficiency be further improved, by punishments of stochastic length? No, this is not possible, when loans are very profitable.[23] Assume that $\frac{x^*}{\bar{K}-1+x^*} > \frac{\ell}{1+\ell}$, so that the condition required for the efficiency result in Proposition 5 is violated. Then there cannot be an equilibrium with a monotone random information structure, since the borrower with credit history B would repay with probability at least $\frac{x^*}{\bar{K}-1+x^*}$. Let the memory length be $K \geq \bar{K}$, and consider the incentives of a borrower who has a single default in the last K periods and this default has happened exactly $\bar{K}$ periods ago. Suppose that the signal structure is random and this borrower is assigned a signal B with probability $x \in (x^*, 1)$ and signal G with complementary probability. Now, if the borrower has a strict incentive to repay after G, she also has a strict incentive to repay after B.[24] Thus under such an information structure, a lender assigns probability at least $\frac{x^*}{\bar{K}-1+x^*}$ to a lender with a bad signal repaying, and will therefore have an incentive to lend.

Non-monotonicity of the information is an unattractive feature, since a borrower with a worse default record is implicitly given a better rating than one with a single default.

[23]Note the contrast with Proposition 5, where the payoff V^* was achievable provided that loans not too profitable.

[24]Indeed, in any purifiable equilibrium, the borrower must play in the same way if she is given a loan, regardless of the signal, since the signal does not affect the history of outcomes.

Furthermore, it may be vulnerable to manipulation, if we take into account real world considerations that are not explicitly modelled. A borrower with two defaults in the last $K-1$ periods is ensured of a continuation value corresponding to being able to default without consequence twice in every K periods. This is very attractive, and a borrower with a single default may be willing to pay a large bribe to a lender, in exchange for the privilege of defaulting once more. Having demonstrated the efficiency results that are feasible when these considerations are absent, we restrict attention to monotone information structures for the remainder of the paper.

3.3 The Irrelevance of Information on Lenders

We have assumed that borrowers cannot observe any information on the past behavior of the lender. Can information on lender behavior be used to prevent lenders from making loans to borrowers who have recently defaulted? Suppose that the borrower observes information on the past outcomes of the lender, just as the lender observes information on the past outcomes of the borrower. It is easy to see that this additional information is irrelevant, and cannot improve outcomes, since observing the outcome in any interaction of the lender does not convey information on whether the lender should have made the loan or not in the first place.

Now, augment the borrower's information, so that she not only observes the outcomes in the lender's past interactions, but also the information that the lender received about the borrowers he interacted with. For example, if the lender's information regarding a borrower is a binary signal, then the borrower today observes the outcome in each past K interaction plus the K realizations of the binary signal observed by the lender. In particular, if the lender lent yesterday to a borrower with a bad signal, the borrower today can see this. The question is: can we leverage this additional information in order to discipline the lender, so that he does not lend to borrowers with a bad signal?

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

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

This discussion also illuminates the interplay between information and incentives that underlies our analysis. Since a lender's future continuation value cannot depend on his current behavior, incentives have to be provided *within the period*. This is possible, since borrowers with a bad signal have a higher incentive to default than those with a good signal. Furthermore, the underlying moral hazard problem is one-sided — the lender has no incentive to deviate if he expects the efficient outcome (i.e. the loan plus repayment). Thus our analysis can be generalized to games where only one player has an incentive to deviate from the path to an efficient outcome, but not when both players have an incentive to deviate. This is examined further in Section 6.

3.4 The Length of Memory

Suppose that the bound on memory, K, is exogenously given. Suppose that $K < \bar{K}$. The following random signal structure can deliver punishments of the correct length, and supports the pure strategy equilibrium where the lender makes a loan if and only if the signal is G. Signal G results if there is no instance of D in the record, and if the most recent record is R, so that the borrower got a loan and repaid in the most recent period. If there is any instance of D in the last K periods, signal B results for sure. If there is no instance of D in the record, and the most recent record is N, so that the borrower was excluded, then signal B results with probability x, and G arises with complementary probability. When $x = 0$, the punishment is too mild, and the borrower has no incentive to repay a loan, while if $x = 1$, the punishment is infinitely long, so that the borrower's incentive constraint (1) is more than satisfied. Thus there exists a value of x such that the borrower's incentive constraint holds, and where her payoff is close to V^*. That value decreases when the length K of the memory increases.

It remains to examine the lender's incentive under this information structure. For a borrower with any instance of D in her record, let $n \in \{1, 2.., K\}$ denote the number of periods that must elapse before the most recent instance of D is eliminated. Let $n = 0$ denote a borrower with signal B but with no instance of D in the last K periods. Finally, let $n = G$ denote a borrower with signal G. Under the equilibrium profile, the invariant distribution over values of n has $\mu_K = \mu_{K-1} = ... = \mu_1$, and $\mu_0 = \frac{x}{1-x}\mu_1$.[25] A borrower with $n \in \{0,1\}$ will repay on being given a loan, while those with $n > 1$ will default. Therefore, using Bayes rule, the probability the lender assigns to a borrower with signal B repaying equals

$$\frac{\mu_0}{\mu_0 + K\mu_1} = \frac{x}{x + (1-x)K}.$$

Notice that the repayment probability is increasing in x, and tends to 1 as $x \to 1$. We have seen that for shorter memories, a larger x is required in order to make punishments severe enough to overcome the borrower's moral hazard. However, this aggravates the lender's moral hazard, since lending to borrowers with a bad signal becomes more profitable. In other words, short memory imposes a cost, and may make it impossible to sustain an efficient equilibrium.

On the other hand, if memory is longer than required, so that $K > \bar{K}$, it is easy to shorten the length of effective memory. An obvious way is to construct an information partition (such as the simple partition in our previous sections) where the partitioning does not depend upon the last $K - \bar{K}$ periods. A second way is to release *complete information* on the outcomes in periods $\bar{K}+1, .., K$. That is, we could construct an information partition with $2 \times |O|^{K-\bar{K}}$ elements, which indicates whether a default took place in the most recent $\bar{K}$ periods or not, and reveals the exact outcome in each period after that. The induction arguments that underlie Proposition 2 can be used to show that in any equilibrium, there cannot be any conditioning on the outcomes in periods $\bar{K}+1, .., K$.

3.5 Infinite Memory Revisited

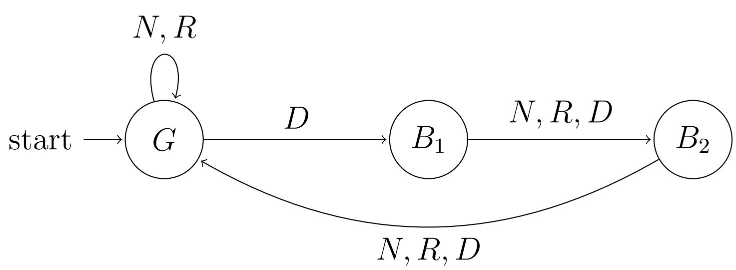

It is useful at this point to examine lender moral hazard when information is perfect and unbounded. We want to ensure that a borrower who defaults, and who should be excluded for K periods, is not offered a loan. To do this, we must distinguish between defaults that

[25]The invariant distribution μ over {G,0,1,2,..,K} has $\mu_G = \frac{(1-x)(1-\lambda)}{x\lambda+(1-x)(1-\lambda+K\lambda)}$, $\mu_0 = \frac{x\lambda}{x\lambda+(1-x)(1-\lambda+K\lambda)}$ and $\mu_1 = \cdots = \mu_K = \frac{(1-x)\lambda}{x\lambda+(1-x)(1-\lambda+K\lambda)}$.

occur when a loan should be made, and those that arise when the lender should not have lent in the first place. This is illustrated in the equilibrium described by the automaton above, where a defaulting borrower is excluded for K periods — the figure depicts the case of $K = 2$. Depending on the entire history, the borrower is either in a good state or in one of K distinct bad states. The lender extends a loan if and only if the borrower is in the good state. A borrower begins in the good state, and stays there unless she defaults, in which case she transits to the first of the bad states. The borrower then transits through the remaining $K - 1$ bad states, spending only one period in each, and then back to the good state. The transition out of any bad state is *independent of the outcome in that period.* This ensuring that the borrower's actions at is in a bad state do not affect her continuation value. Since the borrower is never punished for a default when she is in a bad state, she will always choose to default. This ensures that no lender will lend to her when she is in a bad state. Note that this equilibrium requires that every lender should be able to observe the entire history of every borrower he is matched with. Otherwise, he cannot deduce whether the borrower defaulted in a period where she was supposed to be lent to, or one in which she was supposed to be excluded. [26]

It is interesting at this point to ask, what would happen in an infinite memory context if lenders adopted a simple norm of excluding a defaulting borrower for K periods, independent of the circumstances under which default took place. The unravelling argument would apply in all its force. A lender who observes that a default happened exactly K periods ago who have an incentive to deviate from this norm, reasoning that the borrower has every incentive to repay. In consequence, over time, punishments would be shortened so that the norm became one of exclusion for $K - 1$ periods. Once established, this would induce an enterprising lender to also lend to a borrower who defaulted $K - 1$ periods ago, thereby shortening exclusion further until the length of punishments was insufficient to discipline borrowers.

3.6 An Alternative Specification of Forgetting

We now consider an alternative modelling of forgetting defaults and show that it yields qualitatively similar conclusions. Suppose that the borrower's history is edited, so that any incidence of D is replaced by N after K periods have elapsed, i.e. it is as though the loan

[26] With a public randomization device that is realized at the end of every period, efficiency can be sustained via simpler strategies — a default can be punished by permanent exclusion with some probability. However, the realizations of the randomization device need to be part of the infinite public record.

never took place. Let us also assume that in each period there is a small probability ρ that a borrower does not meets a lender. When a lender and a borrower are matched, the lender perfectly observes the borrower's entire *edited* history, and that $K \geq \bar{K}$, so that an incidence of default is retained longer in the record than is required for incentivizing the borrower.

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

We define the following equivalence relation on t-period histories. Consider two histories, $h^t = (a_1, ...a_t)$ and $\hat{h}^t = (\hat{a}_1, ..., \hat{a}_t)$. We write that $h^t \sim \hat{h}^t$ if for every $s \in \{1, \ldots, t-K\}$, $\hat{a}_s = a_s$, while for every $s \in \{t-K+1, \ldots, t\}$, $\hat{a}_s \neq a_s \Rightarrow \hat{a}_s, a_s \in \{D, N\}$. That is, two t-period histories are equivalent if:

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

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

Lemma 9 *If σ is a purifiable equilibrium, then at every date $t+1$, σ is measurable with respect to H^t, the set of possible recorded histories. Further, if $h^t \sim \hat{h}^t$, $\sigma(h^t) = \sigma(\hat{h}^t)$.*

Proof. See appendix. ■

This lemma implies that in any purifiable equilibrium, not obtaining a loan must be treated in the same way as a default. This has important implications. Consider an equilibrium where borrowers are incentivized to repay by $\bar{K}$ periods of exclusion. The above lemma implies that if any period, a borrower fails to get a loan, then she must also be excluded for $\bar{K}$ periods. Thus if $\rho > 0$, so that there is some chance that a borrower might fail to get a loan for exogenous reasons, the fact that such failures must lead to punishment causes additional inefficiencies. Providing the lender with coarse information can improve efficiency in this context as well. In particular, if the borrower is only informed about the last $\bar{K}$ outcomes, and only learns whether the borrower has defaulted or not in this time, then an equilibrium can be sustained under exactly the same conditions as in Section 3.1.

4 Random Exclusion

Let the memory length be $\bar{K}$, and assume the *simple* information structure, where B is observed by the lender if the borrower defaulted in any of the last $\bar{K}$ periods, and G is observed otherwise. Throughout this section we assume that lending is so profitable that a pure strategy equilibrium does not exist, i.e. $\ell < \frac{1}{\bar{K}-1}$. Suppose that lenders lend with probability $p \in (0,1)$ on observing B, and lend with probability one after G. We now show that this permits an equilibrium where the length of exclusion after a default is no greater than $\bar{K}$ — indeed, the effective length is strictly less, since exclusion is probabilistic. This may appear surprising — if a lender is required to randomize after B, then not lending must be optimal, and so the necessary incentive constraint for an individual lender should be no different from the pure strategy case. However, the behavior of the population of lenders changes the relative proportions of different types of borrower among those with signal B. It raises the proportion of those with larger values of m, thereby raising the default probability at B and disciplining lenders. Thus, other lenders lending probabilistically to borrowers with a bad history exacerbates the adverse selection faced by the individual lender.

Let p denote the probability that a borrower with history B gets a loan (a borrower with history G gets a loan for sure). Recall that if $p = 0$ and $K = \bar{K}$, then it is strictly optimal for a borrower with a good signal, i.e. $m = 0$, to repay. By continuity, repayment is also optimal for a borrower with $m = 0$ for an interval of values, $p \in [0, \tilde{p}]$, where $\tilde{p} > 0$ is the threshold where such a borrower is indifferent between repaying and defaulting. We restrict attention to values of p in this interval in what follows. Note that the best responses of a borrower with $m = 1$ are identical to those of a borrower with $m = 0$, for any p, since their continuation values are identical. Also, any increase in p increases the attractiveness of defaulting, and so a borrower with $m > 1$ will continue to default when $p > 0$.

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

$$\tilde{V}^{\bar{K}}(0,p) = (1-\delta) + \delta \left[\lambda \tilde{V}^{\bar{K}}(\bar{K},p) + (1-\lambda)\tilde{V}^{\bar{K}}(0,p) \right], \tag{12}$$

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

$$\tilde{V}^{\bar{K}}(m,p) = \begin{cases} p(1-\delta) + \delta\left[p\lambda V^{\bar{K}}(\bar{K},p) + (1-p\lambda)V^{\bar{K}}(m-1,p)\right] & \text{if } m = 1, \\ p(1-\delta)(1+g) + \delta\left[pV^{\bar{K}}(\bar{K},p) + (1-p)V^{\bar{K}}(m-1,p)\right] & \text{if } m > 1. \end{cases} \tag{13}$$

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

$$\mu_{m-1} = (1-p)\mu_m \quad \text{if } m > 1. \tag{14}$$

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

$$\mu_{\bar{K}} = \lambda\left(\mu_0 + p\mu_1\right) + p\sum_{m=2}^{\bar{K}} \mu_m. \tag{15}$$

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

$$\mu_0 = (1-\lambda)\mu_0 + (1-p\lambda)\mu_1. \tag{16}$$

Since μ_m depends on p and also on the repayment probability for borrowers with types $m \in \{0,1\}$, which equals 1, we write it henceforth as $\mu_m(p,1)$. Figure 2 depicts the invariant distribution over the values of $m \in \{1,2..,\bar{K}\}$, conditional on signal B, for two values of p. The horizontal line, in red, depicts the conditional distribution when $p = 0$, which is uniform. The distribution conditional on $p > 0$, in blue, is upward sloping since higher values of p increase $\mu_{\bar{K}}$ and depress μ_1.

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

$$\pi(p,1) = \frac{\mu_1(p,1)}{1-\mu_0(p,1)}. \tag{17}$$

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 $\pi(\tilde{p},1) \leq \frac{\ell}{1+\ell}$, the intermediate value theorem implies that there exists a value of $p \in (0,\tilde{p}]$ such that $\pi(p,1) = \frac{\ell}{1+\ell}$. This proves the existence of a mixed

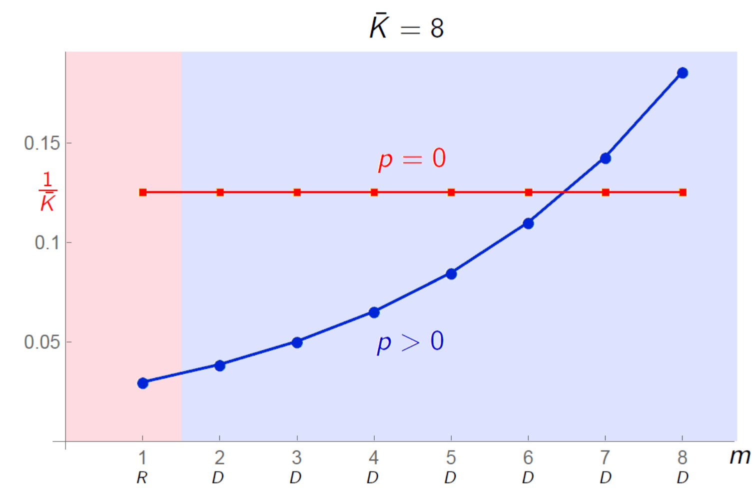

Figure 2: *Stationary probabilities conditional on signal B: $\mu_m/(1-\mu_0)$ for $m = 1, \ldots, \bar{K}$. Illustrated for $p = 0$ and $p = 0.23$.*

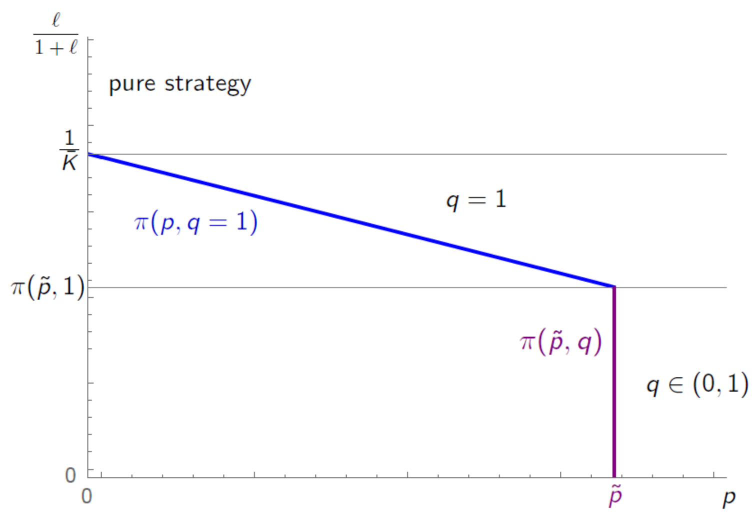

Figure 3: *The functions $\pi(p, 1)$ and $\pi(\tilde{p}, q)$. When $\pi(\tilde{p}, 1) \leq \ell/(1+\ell) \leq 1/\bar{K}$ there exists a mixed strategy equilibrium in which all types of borrowers use a pure strategy. When $0 < \ell/(1+\ell) < \pi(\tilde{p}, 1)$, there exists a mixed strategy equilibrium in which types $m = 0$ and $m = 1$ repay with probability q setting $\pi(\tilde{p}, q) = \ell/(1+\ell)$.*

strategy equilibrium where all borrowers have pure best responses.

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

Let $\mu(\tilde{p}, q)$ denote the invariant distribution over values of m induced by this strategy profile. The probability that a loan made at history B is repaid is now

$$\pi(\tilde{p}, q) = \frac{q\mu_1(\tilde{p}, q)}{1 - \mu_0(\tilde{p}, q)} = q\,\pi(\tilde{p}, 1). \tag{18}$$

We establish the second equality in Appendix A.5.2. Clearly, $\pi(\tilde{p}, q)$ is a continuous and strictly increasing function of q. Since we are considering the case where $\pi(\tilde{p}, 1) > \frac{\ell}{1+\ell}$, and since $\pi(\tilde{p}, 0) = 0$, the intermediate value theorem implies that there exists a value of q setting the repayment probability $\pi(\tilde{p}, q)$ equal to $\frac{\ell}{1+\ell}$, so that lenders are indifferent between lending and not lending to a borrower with signal B. Figure 3 illustrates our analysis, with p on the horizontal axis and the repayment probability after B, $\pi(p, q)$, on the vertical axis. When $q = 1$, so that a borrower with $m = 1$ repays for sure, $\pi(p, 1)$ is given by the downward sloping blue line. This stops at $\tilde{p}$, and further declines of the repayment probability are achieved by reducing q, along the vertical purple line. We have therefore established the following proposition.

Proposition 10 *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 a borrower with signal G is given a loan for sure, and the borrower with signal B is given a loan with positive probability. The borrower's payoff at G in such an equilibrium is strictly greater than in any pure equilibrium under a simple information structure. If ℓ is strictly greater than a threshold value ℓ^*, then $p < \tilde{p}$ and the borrower has strict best responses for every value of m, and 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 L.*

Appendix A.6 proves that the equilibria described in the proposition are purifiable. In the case where both lenders and borrowers with $m \in \{0, 1\}$ mix, it is worth noting that the equilibrium is not *regular*, since it lies in a one-dimensional manifold of equilibria.[27] When the lending probability after B equals $\tilde{p}$, and borrower types $m \in \{0, 1\}$ are indifferent between

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

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 $\pi(\tilde{p}, q) = \frac{\ell}{1+\ell}$, but the repayment probability at $m = 0$ can take any value in some interval. However, the equilibria where the repayment probabilities differ must condition on payoff-irrelevant histories, and the payoff perturbations eliminate such conditioning. Thus, only one of these equilibria can be purified, namely the one where the repayment probability at $m = 0$ also equals $\tilde{q}$. This is true even though there exists a more efficient equilibrium where the borrower repays for sure when $m = 1$.

Borrowers are strictly better off in the mixed equilibrium than in a pure equilibrium with a K strictly larger than $\bar{K}$. To see this, observe that a borrower with a bad signal could always follow a strategy of never defaulting, in which case then she would be excluded only partially (since p is interior) and also for a shorter period of time. The welfare comparison for a lender is less straightforward. A lender always gets an expected payoff of 0 when he is matched with a borrower with signal B. In any equilibrium where borrowers with signal G repay for sure, the lender's expected payoff equals the measure of such borrowers, μ_0. A pure strategy equilibrium with larger K results in a lower μ_0 since exclusion is longer. However, the mixed strategy equilibrium also results in a lower μ_0 since borrowers re-offend. In our numerical examples, μ_0 is larger in the mixed equilibrium than in the pure equilibrium with longer memory. However, lenders can be considerably worse off in the equilibria where borrowers also randomize. In this case, μ_0 is even smaller, in addition to which, even borrowers with a good signal sometimes default willfully.

An Example We consider parameter values such that $\bar{K} = 4$.[28]. Consider the pure strategy profile when $K = \bar{K}$, where the lender extends a loan only after G. From our previous analysis, the invariant distribution over $\{1, \ldots, K\}$ is uniform, and the probability that the lender is repaid if lending at B equals $\frac{1}{4}$. So if $\ell > \frac{1}{3}$, a pure strategy equilibrium exists. The expected payoff to a borrower with a good history is $\bar{V} = 0.763$. The lender's expected payoff equals the probability of encountering a borrower with a good history, $\mu_0 = 0.714$.

If $\ell < \frac{1}{3}$, lending after B is too profitable and a pure strategy equilibrium with $\bar{K}$-period memory does not exist. A pure strategy equilibrium with longer memory exists, but can be very inefficient. For example, if $\ell = 0.315$, we need $K = 30$, in which case $\lfloor m^*(K) \rfloor = 7$ and the lender strictly prefers not lending at B. The invariant distribution over borrower

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

types has only a quarter of the population with a clean history, so that the lender's payoff equals 0.25. Since exclusion is very long, the borrower's payoff at a clean history is also low: $V^K(0) = 0.537$.

In the mixed strategy equilibrium with 4-period memory, the lender offers a loan with probability $\tilde{p} = 0.028$ to a borrower on observing B. This equilibrium is considerably more efficient than the pure equilibrium with 30-period memory. The proportion of borrowers with clean histories is $\mu_0 = 0.705$, and this is also the lender's expected payoff. The expected payoff to a borrower with a clean history is $V^{\bar{K}}(0, \tilde{p}) = 0.775$, which is higher than her payoff under the pure strategy profile with 4-period memory, since she is now sometimes extended 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 $m \in \{0, 1\}$.[29] 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 is $V^{\bar{K}}(0, \tilde{p}) = 0.778$.

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 0.778. 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.

It may be illuminating to examine the findings of this section in the context of the recent literature on information design following Kamenica and Gentzkow (2011). Bergemann and Morris (2016) is particularly relevant since they consider a Bayesian game and show that the information designer optimally provides players with coarse information regarding the state and regarding the (recommended) actions of other players. In our context, the distribution over states is endogenously determined, by the recommended distribution over actions. This is why random actions play a role, since they affect the long run distribution over states.

[29] If $\ell = 0.1$, $K = 89$ periods of exclusion are needed to support a pure strategy equilibrium. In this case $\lfloor m^*(K) \rfloor = 8$. The lender's payoff is $\mu_0 = 0.1$ and the borrower's payoff at a clean history is $V^K(0) = 0.526$.

5 No Loan Chasing

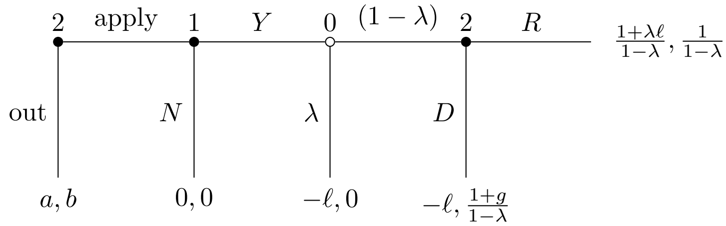

We now investigate the effects of laws mandating that lenders cannot "chase" borrowers, requiring instead that borrowers must first make an application to the lender. Making an application involves a slight cost to the borrower. [30] We model the interaction between borrower and lender as follows. First, the borrower may apply for a loan, at a small cost. If she does not apply, the game ends, with payoffs a for the lender and $b > 0$ for the borrower. 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 backwards induction outcome is *out*, and the borrower does not apply. Assume the simple information structure, under which lender knows only whether the borrower has defaulted in the last K periods (credit history G) or not (credit history B).

Requiring prior application transforms the interaction between an individual lender and borrower into a signaling game. In particular, 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. Thus, in the context of an equilibrium where borrowers with different credit histories are treated differently, the application decision can signal the borrowers's private information.

Proposition 11 *Let* $K \geq \max\{\bar{K}, 2\}$*, and assume the simple information structure. For any* $\ell > 0$*, there exists a sequentially strict 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 ($a > 0$). 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.

[30] And possibly to the lender, although this plays no role in the analysis.

Suppose that a borrower of type $m \geq 1$ with a bad credit history gets a loan upon applying. As in Section 3.1, her repayment incentives are such that she 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

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

for some $m < m^*$.

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, her net benefit from applying equals

$$q\big[(1-\delta)(1+g)+\delta\big(V(K)-V(m'-1)\big)\big]-(1-\delta)a. \tag{20}$$

We now show that expression (20) above is strictly positive. Substituting for $(1-\delta)a$ from (19), we see that the sign of (20) is the same as that of

$$(1-\delta)g+\delta(1-\lambda)\big(V(K)-V(m'-1)\big)+\delta\lambda\big(V(m-1)-V(m'-1)\big). \tag{21}$$

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 (21) is strictly positive. Also, $V(m) > V(m')$ since $m < m'$, and so the third term as well as the overall expression in (21) 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 (Cho and Kreps (1987)) implies that in the above equilibrium, the lender must assign probability one to defaulting types when he sees an application from a borrower with a B credit history. ■

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 $\bar{K}$, as in Section 4, 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 Section 4.

6 Generalizing our Results

The analysis of this paper applies to many large markets with one-sided moral hazard. This includes, for example, the interaction between a buyer and a seller when the seller decides which quality to provide after the buyer places the order. If information about the past outcomes of the seller is subject to a bounded memory, our analysis applies. More generally, our analysis applies to a large class of stage games of perfect information where moral hazard is *effectively* one-sided. We show in the appendix that this includes *any stage game* in which each player moves at most once.

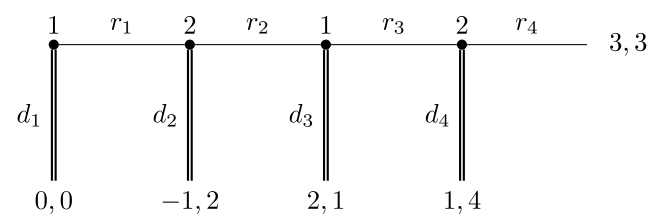

Figure 4: *The centipede game.*

To illustrate how our results generalize, consider the centipede game in Figure 4, where the backwards induction strategy profile has players choosing *down* (d_t) at every history. Consider the Pareto-efficient terminal node z^* that is reached when players choose *right* (r_t) at every history. *If players expect the outcome* z^*, then only one player has an incentive to deviate along the path to z^*: player 2. Furthermore, there is only a single decision node at which 2 has an incentive to deviate: the last one.

Now suppose that this game is played repeatedly in a random matching environment. Consider an information structure where, before the game is played in a match, player 1 is provided coarse information on the last K outcomes of player 2. Specifically, signal B is observed if and only if 2 has played d_4 in at least one of the last K periods; signal G is observed otherwise. Suppose that 1 plays d_1 upon observing signal B, and let $\bar{K}$ denote the smallest punishment length that suffices to incentivize player 2 to play r_4 at G.[31] As

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

long as $K \geq \max\{\bar{K}, 2\}$, there exists an equilibrium that supports the efficient outcome z^* under this information structure. Furthermore, it is costly for player 2 to choose r_2 at her first decision node if she expects player 1 to continue with his backwards induction action d_3. This ensures that if 2 plays r_2 and has a bad signal, then the D1 criterion implies that, at his second information set, 1 believes that if he plays r_3, 2 will continue with d_4.

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

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

Under these assumptions, there exists a simple information structure and an equilibrium that sustains play of the efficient outcome z^*.[32] In this equilibrium, players play the outcome z^* when the signal about player $\hat{i}$ is good, and they play the backwards induction outcome when the signal is bad. Players' beliefs satisfy the D1 criterion, and it suffices to pick punishments of a minimum length that is sufficient to discipline player $\hat{i}$.

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

If the first of these assumptions is not satisfied, and both players have an incentive to deviate from the path to z^*, then providing incentives for both players becomes more difficult. This would be the case, for example, if the payoff of player 1 after r_3 was changed to 4. As we have noted in Section 3.3, it is difficult to provide incentives for a player to condition his behavior on a signal regarding his opponent via *future play.* Incentives for such conditioning have to be provided within the period. (In the lender-borrower game, the lender has only within-period incentives to condition his behavior on the borrower's signal). This makes the analysis of this case more difficult, and we leave this for future work.

[32]In fact, our analysis applies to stage games of perfect information with an arbitrary finite number of players, provided that these conditions hold, but we prove this only for the case of two players.

7 Conclusion

Excessive lending by financial institutions has been the focus of attention during the recent financial and sub-prime crises. Economists and policy-makers (e.g. Admati and Hellwig (2014)) have argued that government guarantees for banks, whether implicit or explicit, induce lending to high-risk borrowers, the cost of which is borne by the government. This paper identifies a different form of excessive lending, whereby lenders undermine borrower repayment incentives by providing defaulters with credit. The negative externality imposed by such behavior falls upon other lenders, and ultimately upon the credit system as a whole.

While this problem is quite general, our specific focus has been on the question: how should information be optimally provided when records are bounded? Providing incentives for a lender to punish a defaulting borrower, by excluding her, turns out to be critical. Even though moral hazard is assumed to be one-sided, a system which provides sufficient incentives for borrowers with clean records to repay also provides similar incentives for borrowers whose record is on the verge of becoming clean. Consequently, lenders have little incentive to exclude them. Our key finding is that, in order to support lending, information needs to be coarse, so that a lender cannot accurately identify which recent defaulter is most likely to repay her loan. Nonetheless, even with coarse information, borrower exclusion may need to be inefficiently long, since disciplining lenders is non-trivial. We have examined how random exclusion may increase efficiency even though it may adversely affects lenders. Finally, legislation that prohibits lenders from "chasing borrowers" effectively discourages applications from borrowers with a bad credit history and improves efficiency.

A Appendix

A.1 Proof Lemma 3

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

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

Since both δ^K and $V^K(0)$ are strictly decreasing in K, it suffices to prove this for $K = \bar{K}$. Here, we prove a stronger result that implies the first part of Lemma 3. For $K = \bar{K}$ and for any $m > 1$,

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

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

$$(1-\delta)g > \delta(1-\lambda)\left(V^{\bar{K}}(1) - V^{\bar{K}}(\bar{K})\right) = \delta(1-\lambda)(\delta - \delta^{\bar{K}})V^{\bar{K}}(0). \tag{A.22}$$

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.22), it suffices to show that

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

Since $V^{\bar{K}-1}(0) > V^{\bar{K}}(0)$, the above inequality is proved. We have therefore established (6) and the first part of Lemma 3. To prove the rest of this lemma, we solve for m^*, the (real) value of m that sets (7) equal to zero:

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

Since

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

it follows that

$$\lim_{K\to\infty} m^*(K) = \ln\left[\frac{g(1-\delta(1-\lambda))}{1-\lambda}\right] / \ln\delta =: m^\infty. \tag{A.23}$$

Therefore, $\frac{\lfloor m^*(K)\rfloor}{K} \to 0$ as $K\to\infty$. This completes the proof of Lemma 3.

A.2 Proof of claims in Section 3.1

A.2.1 First Claim

We show that the right hand side of (11) is a strictly increasing function of x. Using the expression in (10), we obtain that the derivative of the right hand side of (11) is

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

which is strictly positive for all $x\in[0,1]$.

A.2.2 Second Claim

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

$$(1-\lambda)(1-x(1-\delta))\left(\delta^{m-1}-\delta^{\bar K}\right)V^{\bar K,x}(0)-(1-\delta)g \tag{A.24}$$

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.24) equal to zero, any increase in x must be compensated by a decrease in m^*. The result follows.

We now show that $m^*(\bar K,x)\in[1,2)$ for every $x\in[x^*(\bar K),1]$. Since (A.24) is strictly decreasing in m, it suffices to show that

$$(1-\lambda)(1-x^*(1-\delta))\left(\delta^{2-1}-\delta^{\bar K}\right)V^{\bar K,x^*}(0)<(1-\delta)g \tag{A.25}$$

and that

$$(1-\lambda)(1-x^*(1-\delta))\left(\delta^{1-1}-\delta^{\bar{K}}\right)V^{\bar{K},x^*}(0)\geq(1-\delta)g. \tag{A.26}$$

Since (A.24) 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 (11) 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.25).

Similarly, since (A.24) 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.26).

A.3 Proof of Lemma 9

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

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

A.4 Non-stationary analysis

Consider the simple binary partition of Section 3.1 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 $m < m^*(K)$ repay when extended a loan. Thus, at any date t, the probability that a loan extended at history B is repaid is

$$\tilde{\pi}(t; K) = \frac{\sum_{m=1}^{\lfloor m^*(K) \rfloor} \tilde{\mu}_m(t)}{1 - \tilde{\mu}_0(t)}. \tag{A.27}$$

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

$$\bar{\pi}(K) := \sup_t \tilde{\pi}(t; K) < \frac{\ell}{1 + \ell}. \tag{A.28}$$

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

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

so that

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

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

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

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

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

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

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

A.5 Proof of claims in Section 4

A.5.1 First Claim

We show that $\pi(p,1)$, defined in (17), is a strictly decreasing function of p. (Continuity for $p \in [0,1]$ is immediate.) For any given $K \geq 2$ and $p \in (0,1)$, the invariant distribution $\mu(p,1)$, is given by equations (14) to (16), together with the condition $\sum_{m=0}^{K} \mu_m = 1$. Solving the above system, we obtain

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

so that

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

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

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

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

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

A.5.2 Second Claim

We now establish the second equality in (18). For any given $K \geq 2$ and $(p,q) \in (0,1)^2$, the invariant distribution $\mu(p,q)$, is given by

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

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

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

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

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

as in (18).

A.6 Purification of the mixed equilibrium in Proposition 10

The perturbed stage game Γ^ϵ is defined as follows. Without loss of generality, it suffices to perturb the payoff to one of two actions of each of the players. Accordingly, we assume that the payoff to the lender from lending is augmented by ϵy, where y is the realization of a random variable that is distributed on a bounded support, without loss of generality $[0, 1]$, with a continuous cumulative distribution function F_Y. The expected payoff to the borrower from willful default is augmented by ϵz, where z is the realization of a random variable that is distributed on $[0, 1]$ with a continuous cumulative distribution function F_Z.

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

We now turn to the mixed equilibria of Proposition 10. The case where only the lender mixes and the borrower has strict best responses is more straightforward. So we consider first the case where both lender and borrower mix. The equilibrium is not *regular*, since it is contained in a one-dimensional manifold of equilibria. Thus we cannot directly invoke, for example, Doraszelski and Escobar (2010), who show that regular Markov perfect equilibria are purifiable in stochastic games.

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

$$\tilde{V}^{\bar{K}}(0, \bar{p}) = (1-\delta)\left(1 + \epsilon \int_{\bar{z}}^{1} (z - \bar{z})\, dF_Z(z)\right) + \delta \left[\lambda \tilde{V}^{\bar{K}}(\bar{K}, \bar{p}) + (1-\lambda)\tilde{V}^{\bar{K}}(0, \bar{p})\right]. \quad \text{(A.29)}$$

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

$$\tilde{V}^{\bar{K}}(1, \bar{p}) = \bar{p}(1-\delta)\left(1 + \epsilon \int_{\bar{z}}^{1} (z - \bar{z})\, dF_Z(z)\right) + \delta \left[\bar{p}\lambda \tilde{V}^{\bar{K}}(\bar{K}, \bar{p}) + (1-\bar{p}\lambda)\tilde{V}^{\bar{K}}(0, \bar{p})\right]. \quad \text{(A.30)}$$

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

$$\tilde{V}^{\bar{K}}(m, \bar{p}) = \bar{p}(1-\delta)\left(1 + g + \epsilon\, \mathbb{E}(z)\right) + \delta \left[\bar{p}\tilde{V}^{\bar{K}}(\bar{K}, \bar{p}) + (1-\bar{p})\tilde{V}^{\bar{K}}(m-1, \bar{p})\right] \quad \text{(A.31)}$$

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

$$(1-\delta)(g + \epsilon\bar{z}) - \delta(1-\lambda)\left(\tilde{V}^{\bar{K}}(0, \bar{p}) - \tilde{V}^{\bar{K}}(\bar{K}, \bar{p})\right) = 0. \quad \text{(A.32)}$$

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

$$\frac{\bar{q}\mu_1(\bar{p}, \bar{q})}{1 - \mu_0(\bar{p}, \bar{q})} - \frac{\ell}{1 + \ell + \epsilon\bar{y}} = 0. \quad \text{(A.33)}$$

When $\epsilon = 0$, equations (A.32) and (A.33) 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.32) and (A.33) at $\epsilon = 0$ and $(p, q) = (\tilde{p}, \tilde{q})$ is non-zero. By the implicit function theorem, if ϵ is sufficiently close to zero, there exist $(\bar{p}(\epsilon), \bar{q}(\epsilon))$ close to $(\tilde{p}, \tilde{q})$ that solves equations (A.32) and (A.33). We have then established that the mixed strategy equilibrium where both the lender and the borrower mix is purifiable.

The indifference condition for a borrower with type $m \in \{0, 1\}$ in the unperturbed game is given by[33]

$$\gamma(p,q) = (1-\delta)g - \delta(1-\lambda)(V^{\bar{K}}(0,p,q) - V^{\bar{K}}(\bar{K},p,q)) = 0. \tag{A.34}$$

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

$$\phi(p,q) = q\mu_1(p,q)(1+\ell) - \ell(1-\mu_0(p,q)) = 0. \tag{A.35}$$

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

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

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

Differentiating with respect to p, we obtain

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

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

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

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

Differentiating ϕ with respect to q gives

$$\phi_q(p,q) = \frac{\partial \mu_1(p,q)}{\partial q} q(1+\ell) + \mu_1(p,q)(1+\ell) + \frac{\partial \mu_0(p,q)}{\partial q}\ell. \tag{A.36}$$

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

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

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

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

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,

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

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

Using these expressions in (A.36), we obtain

$$\phi_q(p,q) = \mu_1(p,q)(1+\ell)\left[1 - \frac{QC}{(1-Q)(A-QB)}\right] + \frac{(1-\lambda)C}{(1-Q)(A-QB)} \left(1-\mu_0(p,q)\right)\ell. \tag{A.37}$$

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

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

Using this indifference condition in (A.37) gives

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

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

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

A.7 General Games

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

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

Definition 12 *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 13 *No player has an incentive to deviate from $\phi(\bar{z})$, the backwards induction path. If $z \neq \bar{z}$, then at least one player has an incentive to deviate from the path to z.*

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

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

The following lemma shows that the set of pairs (Γ, z^*) satisfying the conditions above includes every outcome that Pareto-dominates the backwards induction outcome in *games where each player moves at most once along any path of play.*

Lemma 14 *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^*.*

Proof. Let Γ be a two-player game of perfect information where each player moves at most once along any paths of play. Without loss of generality, player 1 moves at the initial node, choosing an action from a finite set A_1 and player 2 moves after some choices of player 1. Let a_1^*, a_2^* denote the path to z^* from the initial node. We claim that at a_1^*, player 2 has *an incentive to deviate*, i.e. her optimal action differs from a_2^*. If this were not the case, then since player 1 prefers z^* to $\bar{z}$, z^* would be the unique backwards induction outcome, a contradiction. To see that player 1 does not have an incentive to deviate at the initial node x_0, observe that $u_1(\bar{\sigma}(x_0)) = 0 < u_1(z^*)$. ■

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

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

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

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

For instance, in the centipede game depicted in Figure 4, we have $\hat{x} = (r_1, r_2, r_3)$ and $\tilde{x}_j = (r_1, r_2)$.

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

We may therefore define

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

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

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

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

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

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

The repeated game strategies are as follows.

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

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

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

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

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

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

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

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

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

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

Suppose that the pair (Γ, z^*) also satisfies the third assumption set out in Section 6, i.e.

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

In this case it is costly for $\hat{i}$ to continue on the path $\phi(z^*)$ at x, given that j plays the backwards induction strategy at $\tilde{x}_j$. Thus the D1 criterion implies that if the decision node $\tilde{x}_j$ is reached, the probability π assigned by j must be zero. The arguments for this are identical to those set out in the proof of Proposition 11. Thus we have an equilibrium that supports the outcome z^* without any further assumptions.

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

References

Admati, A., and M. Hellwig (2014): *The bankers' new clothes: What's wrong with banking and what to do about it*. Princeton University Press.

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

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

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

Chatterjee, S., D. Corbae, M. Nakajima, and J.-V. Ríos-Rull (2007): "A quantitative theory of unsecured consumer credit with risk of default," *Econometrica*, 75(6), 1525–1589.

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, Yale.

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

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.

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

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

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

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

Fay, S., E. Hurst, and M. J. White (2002): "The household bankruptcy decision," *The American Economic Review*, 92(3), 706–718.

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.

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.

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. (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*, 151, 2–29.

Mathevet, L., J. Perego, and I. Taneva (2016): "Information Design: The Random Posterior Approach," .

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

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

Silvia, J. (2008): "Subprime credit: The evolution of a market," *Business Economics*, 43(3), 14–22.

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

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

Taneva, I. A. (2016): "Information design," .

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