

Community Enforcement of Trust with Bounded Memory

V. BHASKAR

University of Texas at Austin

and

CAROLINE THOMAS

University of Texas at Austin

First version received January 2018; Editorial decision August 2018; Accepted August 2018 (Eds.)

We examine how trust is sustained in large societies with random matching, when records of past transgressions are retained for a finite length of time. To incentivize trustworthiness, defaulters should be punished by temporary exclusion. However, it is profitable to trust defaulters who are on the verge of rehabilitation. With perfect bounded information, defaulter exclusion unravels and trust cannot be sustained, in any purifiable equilibrium. A coarse information structure, that pools recent defaulters with those nearing rehabilitation, endogenously generates adverse selection, sustaining punishments. Equilibria where defaulters are trusted with positive probability improve efficiency, by raising the proportion of likely re-offenders in the pool of defaulters.

Key words: Trust game, Repeated games with community enforcement, Imperfect monitoring, Bounded memory, Credit markets, Information design

JEL Codes: C73, D82, G20, L14, L15

1. INTRODUCTION

We examine information and rating systems designed to induce cooperation, in large societies with bilateral interactions and one-sided moral hazard. Our leading application is the trust game, which captures many economic interactions, such as between buyer and seller, or lender and borrower. Since each pair of agents transacts infrequently, opportunistic behaviour (by the seller or borrower) can be deterred only if it results in future exclusion.

We assume that information on past transgressions is subject to *bounded social memory* and is retained only for a finite length of time. While plausible in any context, this is legally mandated in consumer credit markets. In the U.S., the bankruptcy “flag” of an individual filing for bankruptcy under Chapter 7 remains on her record for 10 years, and must then be removed; if she files under Chapter 13, it remains on her record for 7 years. [Elul and Gottardi \(2015\)](#) find that among the 113 countries with credit bureaus, 90% have time-limits on the reporting of adverse information concerning borrowers.¹ Bounded memory also arises under policies used by internet platforms

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

The editor in charge of this paper was Adam Szeidl.

to compute the scores summarising their participants' reputations. For example, Amazon lists a summary statistic of seller performance over the past 12 months. In the U.S., twenty-four states and many municipalities² have introduced "ban the box" legislation, prohibiting employers from asking job applicants about prior convictions unless those relate directly to the job.³

How do societies enforce trustworthiness when constrained by bounded memory? Consider the credit market interpretation of the trust game, and suppose that each borrower–lender pair interacts only once. Lending is efficient and profitable for the lender, provided the borrower intends to repay the loan. However, the borrower has a short-term incentive to wilfully default. Thus, lending can only be supported via long-term repayment incentives whereby default results in the borrower's future exclusion from credit. Since a borrower may sometimes default involuntarily, efficiency requires that exclusion only be temporary.

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

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. This turns out to be false. Perfect information on the recent past behaviour of the borrower, in conjunction with bounded memory, precludes any lending, because it allows lenders to cherry-pick those borrowers with the strongest long-term incentives to repay. The intuition is best illustrated using a candidate pure strategy equilibrium with temporary exclusion, where every player has strict incentives at every information set. 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 repays a loan whenever a borrower with a clean record does. Lenders, who are able to distinguish her from more recent defaulters, find it profitable to extend her a loan, thereby reducing the length of her punishment. Repeating this argument, by induction, no length of punishment can be sustained. As a result, no lending can be supported, because lenders cannot be disciplined to not make loans to borrowers with a bad record.

The result that no lending can be supported extends to any sequentially strict equilibrium. In fact, it extends to all purifiable equilibria, as they are sequentially strict in nearby perturbed games. Suppose that lenders and borrowers are affected by small i.i.d. shocks that alter the lender's opportunity cost of lending and the borrower's benefit from defaulting. When these shocks have a continuous distribution, any equilibrium must be strict, since a player is almost never indifferent between two actions. This excludes, in particular, belief-free type equilibria, where a lender is always indifferent between lending and not lending, and breaks this indifference differently depending upon the borrower's record. Such equilibria disappear in the presence of small payoff shocks, a serious weakness for the analysis of credit markets, where such shocks are a likely feature of the real-world environment.

This negative result leads us to explore information structures that provide the lender with simple, binary information about the borrower's history. Specifically, the lender is told only whether the borrower has ever defaulted in the past K periods (labelled a *bad credit history*) or not (labelled a *good credit history*). A borrower's long-term incentives to repay a new loan differ according to

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

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

the most recent default in her history. More recent defaulters, with most of their exclusion phase ahead of them, have a stronger incentive to recidivate. But since lenders do not have precise information on the timing of defaults, they are unable to target their loans to defaulters who are more likely to repay.

Coarse information therefore generates *endogenous adverse selection* among the pool of borrowers with a bad credit history, thereby mitigating the tendency of the lender to undermine punishments. Our question is: how can coarse information and the consequent adverse selection be tailored to sustain efficient outcomes?

The simple information structure just described prevents a total breakdown of lending. If the punishment phase is sufficiently long, the pool of lenders with a bad credit history is sufficiently likely to re-offend, on average, as to dissuade rogue loans by the lender. But depending on the (exogenous) profitability of loans, the length of exclusion may be longer than is needed to discipline borrowers.

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

Our analysis has direct implications for the study of credit markets, where bankruptcy flags must be removed from borrowers' records after a fixed length of time. Empirical evidence from the U.S. shows that consumers experience a jump in credit scores in the quarter when their flag is removed, leading to a large increase in their credit limits and borrowing.⁴ For example, Dobbie *et al.* find that the increase in credit scores corresponds to an implied 3 percentage point reduced default risk, on a pre-flag-removal risk of 32%. Thus information leaves the market when flags are dropped, and memory constraints are real.

Our theory predicts unravelling: lenders should use their precise information on default dates to target loans to borrowers whose flag is about to disappear. Existing empirical work is silent on unravelling, since it has focused entirely on the comparison before versus after flag removal, and does not examine the dynamic path of lending prior to removal. This is a fruitful area for future empirical research. Our model suggests that, to combat unravelling and support efficient lending, lenders should be provided with coarse information regarding defaults. For instance, a default flag should only reveal that a consumer declared bankruptcy in the past 7 or 10 years, but should *not* provide information on the precise date. Of course, such information is valuable to the lender, and concealing it might conflict with the primary objective of credit rating agencies, which is to serve individual lenders' best interests, rather than sustaining socially efficient outcomes.

If there is large unobserved heterogeneity in a borrower's likelihood of involuntary default, unravelling may be mitigated. However, better information on borrower characteristics reduces unobserved heterogeneity and will increase unravelling. The lender may infer a borrower's propensity to default from her past defaults, but also from any additional information, *e.g.* demographics or consumption choices. With improvements in information-processing and the

4. See Musto (2004), Gross *et al.* (2016) and Dobbie *et al.* (2016).

rise of big data, we would expect this second channel to gain prominence, and for the incremental informativeness of default flags to diminish. Currently, there is no consensus on the latter count: Musto and Dobbie *et al.* find that removed flags are predictive of future credit delinquency, even after conditioning on other available information, in particular a consumer's credit score, while Gross *et al.* find no predictive effect.

The remainder of this section discusses the related literature. Section 2 sets out the model. Section 3 derives the constrained efficient benchmarks, which can be attained with infinite memory. It also shows that with bounded perfect memory, no lending can be supported in any purifiable equilibrium. Section 4 shows that a simple information structure prevents the breakdown of lending. Section 5 shows that such an information structure ensures constrained efficient payoffs for the borrower, either via pure strategies or mixed strategies, and also ensures high lender payoffs if the interaction must be initiated by the borrower. Section 6 discusses extensions and the final section concludes.

1.1. Related literature

Kandori (1992), Ellison (1994) and Deb (2008) study community enforcement in a small population where players are randomly matched in each period to play the prisoner's dilemma. Even if a player has no information on his partner's previous behaviour, contagion strategies can be used to support cooperation.⁵ In large populations, contagion strategies cannot be effective, and so a player must have some information on her opponent's past behaviour.⁶ Takahashi (2010) shows that if each player observes the entire sequence of past actions taken by her opponent, or observes the *action profile* played in the previous period by her opponent and her opponent's partner, cooperation can be supported by using "belief-free" type strategies, where a player is always indifferent between cooperating and defecting.⁷ Heller and Mohlin (2017) assume that a player observes a random sample of the past actions of her opponent, and that a small fraction of players are commitment types—an assumption that enables them to rule out belief-free strategies. In both papers, when only the partner's action in the previous period is observed, grim-trigger strategies sustain cooperation if and only if the prisoner's dilemma game is supermodular.

Our setting is different since moral hazard is one-sided. This feature is immediate in the trust game. But it also arises in *any sequential-move game where each player moves at most once*, such as the sequential-move prisoner's dilemma. One-sided moral hazard arises in many applications: ensuring quality in product markets (Klein and Leffler, 1981); deterring opportunism in bilateral trade (Greif, 1993); and groups that have collective reputations (Tirole, 1996). Deb and González-Díaz (2010) analyse simultaneous-move stage games with one-sided moral hazard that are played in a random matching environment. Karlan *et al.* (2009) examine the role of friendship ties in sustaining lending in a network.

Liu and Skrzypacz (2014) analyse seller reputations when buyers are short-lived and have bounded information about the seller's past decisions. Since the seller has a greater incentive to cheat when sales are larger, equilibria display a cyclical pattern where the seller builds up his

5. Nava and Piccione (2014), Wolitzky (2012) and Ali and Miller (2013) analyse community enforcement where a network determines the interaction structure.

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

7. The two cases are closely related to the belief-free strategies considered in Piccione (2002) and Ely and Välimäki (2002) respectively.




FIGURE 1

Extensive and strategic forms of the Trust Game Γ . (a) extensive form; (b) strategic form.

reputation before milking it.⁸ [Ekmekci \(2011\)](#) studies a reputation model with a long-run player and a sequence of short-run players, and shows that bounded memory allows reputation to persist in the long run, even though it dissipates with unbounded memory.

Our work differs methodologically from recent research on repeated games due to our insistence on purifiable equilibria, as in [Harsanyi \(1973\)](#). This rules out equilibria in belief-free type strategies. While belief-free equilibria play a major role in establishing a folk theorem in repeated games with private monitoring (see [Sugaya \(2013\)](#)), they may be unrealistic, and purification offers a way of making this criticism precise. Our positive results, that efficiency can be sustained with simple strategies, differ from the negative results in [Bhaskar \(1998\)](#) and [Bhaskar et al. \(2013\)](#), where purifiability, in conjunction with bounded memory, results in a total breakdown of cooperation.

Endogenous adverse selection plays a central role in our analysis: although the underlying environment has moral hazard but no adverse selection, an optimal information structure does not fully reveal the borrower's recent history to the lender. This idea has precursors in repeated games with private monitoring ([Sekiguchi, 1997](#); [Bhaskar and Obara, 2002](#)), where a player randomizes over two pure strategies, so that the private signals observed by the opponent are informative of the player's continuation play. Similarly, in [Rahman \(2012\)](#) both the worker and her monitor randomize in order to incentivize each other, to work and to monitor, respectively.

2. THE MODEL

Time is discrete and the horizon infinite. In each period, individuals from a continuum population 1 are randomly matched with individuals from a continuum population 2, to play the trust game Γ illustrated in Figure 1a. Player 1 moves first, choosing whether to trust (Y) player 2 or not (N). If he chooses N , the game ends, and both parties get a payoff of zero. If he chooses Y , then player 2 must decide whether to repay this trust (R), or to default (D). However, with a small probability λ , player 2 is unable to repay trust, and is constrained to default. It is profitable for player 1 to trust player 2 if the latter intends to repay, and unprofitable if she intends to default. Moreover, wilful default is profitable for player 2. The strategic form of the game, given in Figure 1b, clarifies the players' incentives: since $g > 0$ and $\ell > 0$, it is a one-sided prisoner's dilemma. The key features of the trust game are:

- The outcome of the backwards induction profile (N, D), where player 1 chooses N and player 2 chooses D , is Pareto-dominated by the (random) outcome that results when the players play (Y, R).
- If players expect (Y, R) to be played, only player 2 has an incentive to deviate.

8. See also [Sperisen \(2018\)](#).

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

Let Γ^∞ denote the infinitely repeated game where at every period players are randomly matched to play the trust game Γ . The borrower has a discount factor $\delta \in (0, 1)$. The discount factor of the lender is irrelevant for positive analysis.⁹ Since the borrower has a short-term incentive to default, incentives to repay can only be provided by her future lenders. The information that these lenders have about the borrower's past behaviour will be based on her last K outcomes in \mathcal{O} , where $\mathcal{O} = \{N, R, D\}$ is the set of *observable outcomes* in the stage game, composed of the events: no loan, repayment, and default. Involuntary and voluntary defaults cannot be distinguished by anyone other than the borrower concerned, and are both denoted by D .

We consider stationary Perfect Bayesian Equilibria, where agents are sequentially rational at each information set, and their beliefs are given by Bayes' rule wherever possible. We focus on equilibria where all lenders play the same strategy, and all borrowers play the same strategy. We would also like our equilibria to be robust. A strong notion of robustness is sequential strictness:

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

Sequential strictness is a demanding requirement, possibly too demanding, as it rules out any equilibrium in mixed strategies. A weaker criterion is to require sequential strictness in a “nearby game”. We argue that, in reality, both lenders and borrowers face random payoff shocks in each period, that affect the opportunity cost of funds and the benefits of defaulting. Thus, the unperturbed game is an idealization, and sequential strictness in the nearby perturbed game is an adequate robustness criterion. When this is met, the equilibrium is said to be *purifiable*, as in [Harsanyi \(1973\)](#).

The remainder of this section makes the robustness criterion precise. Define $\Gamma(\varepsilon)$, a perturbed version of the stage game Γ , indexed by $\varepsilon > 0$, a scaling parameter. Let X denote the set of player decision nodes in Γ . At each node $x \in X$, the payoff of the player who moves from one of her (two) actions is augmented by εz_x , where z_x is the realization of a random variable Z_x with bounded support. The random variables $\{Z(x)\}_{x \in X}$ are independently distributed, and their distributions are atomless. The player who moves at node x observes the realization z_x before she moves. In the repeated version of the perturbed game, $\Gamma^\infty(\varepsilon)$, we assume that the shocks for any player are independently distributed across periods.¹⁰ In the lender–borrower interpretation of the trust game, the lender gets an idiosyncratic payoff shock from making a loan, while the borrower gets an idiosyncratic shock from wilfully defaulting. An equilibrium σ of Γ^∞ is *purifiable* if there exists a sequence of equilibria $\sigma(\varepsilon_n)$ of $\Gamma^\infty(\varepsilon_n)$ that converges to σ for any strictly positive sequence $\varepsilon_n \rightarrow 0$.

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

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

The following lemma, proved in Section B.1.3 of the Online Appendix, shows that purifiability is a generalization of sequential strictness.

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

3. BENCHMARKS

3.1. Infinite memory

Suppose that each lender can observe the entire history of outcomes in \mathcal{O} of each borrower he is matched with, and that the borrower observes no information about the lender. Suppose payoff parameters are such that there exists an equilibrium where lending takes place.¹¹

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 $(1-\delta)g$.¹² The deviation makes a difference to her continuation value only when she is able to repay, *i.e.* with probability $1-\lambda$. Suppose that after a default, wilful or involuntary, she is excluded from the lending market for K periods. The incentive constraint ensuring that she prefers repaying when able is then

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

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

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

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

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

The equilibrium with \bar{K} periods of exclusion can be improved upon—due to integer constraints, the punishment is strictly greater than what is required to ensure borrower repayment. In Appendix A.1 we show that the highest payoff the borrower can achieve in *any* equilibrium is given by:¹³

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

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

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

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

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

To summarize: \bar{V} and \bar{W} are the *constrained efficient payoffs* for the borrower and lender respectively, that reflect both the integer constraint and the borrower's incentive constraint under imperfect monitoring. V^* and W^* are the *fully efficient payoffs*—these reflect the incentive constraint for the borrower, but no integer constraints. We assume that the designer's objective is to achieve a payoff no less than \bar{V} for the borrower.

3.2. Perfect bounded memory

Henceforth, we shall assume bounded memory: at every stage, the lender observes a bounded history of length K of past outcomes in \mathcal{O} of the borrower he is matched with in that stage. (The borrower does not observe any information regarding the lender.) Our first result is a negative one—if the lender has full information regarding the past K outcomes of the borrower, then no lending can be supported.

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

The proof is an adaptation of the argument in Bhaskar *et al.* (2013). To get some intuition, suppose that $K \geq \bar{K}$ and the information partition is the finest possible, and consider a candidate equilibrium where a borrower is lent to unless her record has any instance of \mathcal{D} in the last K periods. A borrower with a clean record prefers to repay. Now consider a borrower with exactly one default that occurred exactly K periods ago. She has incentives identical to those of a borrower with a clean record, and will also repay. Therefore, a lender has every incentive to lend to such a borrower, undermining her punishment. An induction argument then implies that no length of punishment can be sustained.

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

Having established that neither lender nor borrower can condition their behaviour today on ω^K , we can repeat the same argument to show that a borrower will not condition her behaviour on ω^{K-1} , the outcome $K - 1$ periods ago. An induction argument implies that neither lender nor borrower will condition their behaviour on any of the borrower's history. It follows that the only

sustainable equilibrium outcome corresponds to the backwards-induction profile being played in every period.

The weaker requirement of purification ensures that even if an equilibrium of the unperturbed game is not sequentially strict, it must be sequentially strict in the perturbed game. (A player can be indifferent between two actions only on a measure zero set, and her behaviour on a negligible set is of no consequence to other players.) Thus the argument made in the previous paragraphs applies in *any* perturbed game. As a result, although there are belief-free equilibria in the unperturbed game that sustain lending, these have no counterpart when there are payoff shocks, and are therefore not purifiable. An example is as follows. The lender lends with probability one if the borrower's last outcome is \mathcal{R} , and with probability $p < 1$ if the last outcome is \mathcal{D} or \mathcal{N} , where p is chosen to make the borrower indifferent between repaying and defaulting. The borrower defaults with probability q whenever she gets a loan, independent of her previous history, where q makes the lender indifferent between lending and not lending. Although the lender faces the same default probability independent of the borrower's record, she lends with different probabilities depending on the borrower's record. The problem is that, in real-world credit markets, idiosyncratic shocks affect the opportunity cost of funds for the lender, and he will condition his behaviour on these shocks rather than on the borrower's record, so that the belief-free style equilibrium disappears.

The second part of the proposition, that there can be no conditioning on history if $K = 1$, applies to any information structure. To support lending, we will therefore need $K \geq 2$. The next sections show how coarse information can prevent a breakdown of lending, and achieve efficient outcomes.

4. INFORMATION

We have in mind a designer or social planner who, subject to memory being bounded, designs an information structure for this large society, and recommends a non-cooperative equilibrium to the players.¹⁴ The designer's goal is to achieve a borrower payoff no lower than \bar{V} , and a lender payoff no lower than \bar{W} —our focus is mainly on the former. This requires supporting equilibria where lending is sustained, and where a borrower who defaults is not excluded for longer than necessary. At a general level, ours is an instance of information design (*e.g.* [Kamenica and Gentzkow, 2011](#)) although our methods are very different from the approach taken in this literature.

Let K denote the bound on memory chosen by the designer—we allow K to be arbitrarily large but finite. An information system provides information to the lender based on the past K outcomes in \mathcal{O} of the borrower. We assume that the borrower does not receive information on the past outcomes of the lender.¹⁵ Information structures fall into two broad categories. A *deterministic* (or signal) structure consists of a finite signal space S and a mapping $\tau : \mathcal{O}^K \rightarrow S$. More simply, it consists of a partition of the set of K -period histories, \mathcal{O}^K , with each element of the partition being associated with a distinct signal in S , and can also be called a *partitional* information structure. A *random* (or signal) structure allows the range of the mapping to be the set of probability distributions over signals, so that $\tau : \mathcal{O}^K \rightarrow \Delta(S)$.

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

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

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

4.1. *A simple information structure*

Since the length of memory, K , can be chosen without constraints, we assume that $K \geq \max\{\bar{K}, 2\}$.¹⁶ The information structure is given by the following binary partition of \mathcal{O}^K , which we call the *simple* information structure. The lender observes a “bad credit history”, signal B , if and only if the borrower has had an outcome of \mathcal{D} in the last K periods, and observes a “good credit history”, signal G , otherwise. This section shows that lending can be supported under the simple information structure.

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

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

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

Since $K \geq \bar{K}$, the borrower strictly prefers to repay at a good credit history. Let us examine the borrower’s repayment incentives when the lender sees a bad credit history. Note that this is an unreached information set at the candidate strategy profile, since the lender is making a loan when he should not. Repayment incentives are summarized by m . The borrower’s incentives at $m=1$ are identical to those at $m=0$ —for both types of borrower, their current action has identical effects on their future signal. Therefore, a borrower of type $m=1$ will always repay.

Now consider the incentives of a borrower with $m=K$. By repaying, she shortens her punishment length by one period, to $K-1$. Thus default is optimal if

$$(1-\delta)g > \delta(1-\lambda)[V^K(K-1) - V^K(K)] = (1-\lambda)(1-\delta)\delta^K V^K(0). \quad (3)$$

Since δ^K and $V^K(0)$ are strictly decreasing in K , so is their product, which converges to zero as $K \rightarrow \infty$. Consequently, there exists a smallest $\tilde{K} \geq \bar{K}$ such that (3) is satisfied for all $K \geq \tilde{K}$. Appendix A.4 shows that $\tilde{K} = \bar{K}$ when $\bar{K} > 1$. If $\bar{K} = 1$, then $\tilde{K} > \bar{K}$, since type $m=1$ always repays. Assume henceforth that $K \geq \tilde{K}$.

Finally, consider the incentives to repay for a borrower with an arbitrary m . By repaying, the length of exclusion is reduced to $m-1$, while by defaulting, it increases to K . The difference between the value from defaulting and the value of repaying equals

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

16. If actual memory K' is greater than K , it is straightforward to reduce its effective length to K by not disclosing any information about events that occurred more than K periods ago. More subtly, this can also be achieved by *full disclosure* of events that occurred between K and K' periods ago—this follows from arguments similar to those underlying Proposition 1.

The right-hand side of the above expression is defined for all real-valued m , and we have established that it is positive at $m=K$ and negative at $m=1$. Thus there exists a real number, denoted $m^\dagger(K) \in (1, K)$, that sets the payoff difference equal to zero. For generic payoffs, $m^\dagger(K)$ is not an integer, and we assume this to be the case, ensuring that borrowers have strict incentives for each value of m , and that the equilibrium is sequentially strict. Let $m^*(K) := \lfloor m^\dagger(K) \rfloor$ denote the integer value of m^\dagger . If $m > m^*(K)$, the borrower strictly prefers D when offered a loan. If $m \leq m^*(K)$, she strictly prefers R . Intuitively, a borrower who is close to getting a clean history will not default, just as a convict nearing the end of her sentence will be on her best behaviour.

Since the lender has imperfect information regarding the borrower's K -period history, we compute the lender's beliefs about those histories using Bayes' rule. We focus on lender beliefs in the steady state, *i.e.* under the invariant distribution over a borrower's private histories induced by the strategy profile. (Section B.3 of the Online Appendix gives the conditions under which our strategies are optimal in the initial periods of the game, when the distribution over borrower types may differ from the stationary one.) In every period, the probability of involuntary default is constant and equals λ . Under our strategy profile, a borrower with a bad credit history never gets a loan and hence transits deterministically through the states $m=K, K-1, \dots, 1$. Therefore, the induced invariant distribution over values of m gives equal probability to each of these states. Consequently, the lender attributes probability $\frac{m^*(K)}{K}$ to a borrower with signal B repaying a loan. Simple algebra shows that lending to a borrower with a bad credit history is strictly unprofitable for the lender if

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

Suppose that ℓ is large enough that the lender's incentive constraint (5) is satisfied. Then, he finds it strictly optimal not to lend after B , and to lend after G . Thus, there exists an equilibrium that is sequentially strict (and therefore purifiable). In other words, providing the borrower with coarse information, so that he does not observe the exact timing of the most recent default, overcomes the impossibility result in Proposition 1. Even though those types of borrowers who are close to "getting out of jail" would choose to repay a loan, the lender is unable to distinguish them from those whose sentence is far from complete. He therefore cannot target loans to the former.

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

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

The simple information structure provides the lender with coarse information about the borrower's outcomes, generating uncertainty about the borrower's private history and preventing the lender from cherry-picking among borrowers with a bad credit history. Giving the lender coarse information pools his incentive constraints, so they only need to hold on average. With fine information, the lender's incentive constraint may be violated just for one type, but this suffices to cause unravelling and a total breakdown, as in Proposition 1. In other words, coarse information endogenously generates borrower adverse selection, which disciplines the lender.¹⁷

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

5. EFFICIENT EQUILIBRIA

In this section, we investigate the conditions under which an equilibrium with punishments of minimal length (\bar{K}) exists, for $\bar{K} \geq 2$, under the simple information structure. We show that the borrower's constrained efficient payoff \bar{V} can be achieved for all parameter values. When lending is very profitable (ℓ is small), the equilibrium has lenders making loans with positive probability to bad credit risks, resulting in low profits for lenders.

5.1. *Pure strategy equilibrium when ℓ is large*

Consider the pure strategy profile set out in the previous section, with \bar{K} memory. When punishments are of the minimal length \bar{K} , any punishment whose effective length is shorter will not incentivize repayment. Consequently, $m^*(\bar{K}) = 1$ (see Appendix A.4), so that every type $m > 1$ defaults, giving rise to a steady-state repayment probability of $\frac{1}{\bar{K}}$. The lender has strict incentives not to lend to a borrower with a bad credit history if $\frac{1}{\bar{K}} < \frac{\ell}{1+\ell}$, or, equivalently,

$$\ell > \frac{1}{\bar{K} - 1}.$$

Given that punishments are of length \bar{K} , a borrower with a good credit history has a strict incentive to repay. Thus we have a sequentially strict equilibrium that achieves the payoff \bar{V} for the borrower and \bar{W} for the lender. We can achieve the fully efficient payoffs V^* and W^* by using a random signal structure. Define the *random version of the simple information structure* as follows. If there is no instance of \mathcal{D} in the last \bar{K} periods, signal G is observed by the lender. If there is any instance of \mathcal{D} in the last $\bar{K} - 1$ periods, then signal B is observed. Finally, if there is a single instance of \mathcal{D} in the last \bar{K} periods and this occurred exactly \bar{K} periods ago, signal B is observed with probability x and G with probability $1 - x$. Let $x > x^*$, where x^* denotes the value where the borrower is indifferent between repaying and defaulting when she has signal G . In Appendix A.5 we show that $m^* = 1$ under this random signal structure, so that the repayment probability at signal B remains low enough that lending is not profitable, thereby proving the following proposition:

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

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

5.2. *Mixed strategy equilibrium when ℓ is small*

Consider now the case where $\ell < \frac{1}{\bar{K} - 1}$, and suppose that lenders lend with positive probability upon observing B . This permits an equilibrium where the length of exclusion after a default is

no greater than \bar{K} —the effective length is strictly less, since exclusion is probabilistic. This may appear surprising—if a lender is required to 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 behaviour of the population of lenders changes the mix of different types of borrower among those with signal B , raising the proportion of those with larger values of m . This raises the default probability of borrowers with a bad credit history, and disciplines lenders.

Consider the strategy profile with \bar{K} memory where the lender always extends a loan at signal G , and with probability p at signal B , and a borrower with $m > 1$ never repays the loan, and repays with probability q if $m = 0$ or $m = 1$. Recall that if $p = 0$ and $K = \bar{K}$, then it is strictly optimal for a borrower with a good signal, *i.e.* $m = 0$, to repay. By continuity, repayment is also optimal for a borrower with $m = 0$ for an interval of values, $p \in [0, \tilde{p}]$, where $\tilde{p} > 0$ is the threshold where she is indifferent between repaying and defaulting. The best responses of a borrower with $m = 1$ are identical to those of a borrower with $m = 0$, for any p , since their continuation values are identical. Also, any increase in p increases the appeal of defaulting, so a borrower with $m > 1$ will continue to default when $p > 0$.

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

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

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

In Section B.1.1 of the Online Appendix, 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




FIGURE 2

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

the conditional distribution. Thus, if $\pi(\tilde{p}, 1) \leq \frac{\ell}{1+\ell}$, then there exists a value of $p \in (0, \tilde{p}]$ such that $\pi(p, 1) = \frac{\ell}{1+\ell}$. This proves the existence of a mixed strategy equilibrium where all borrowers have pure best responses.

If loans are so profitable that $\pi(\tilde{p}, 1) > \frac{\ell}{1+\ell}$, then an equilibrium also requires mixing by the borrower. At \tilde{p} , the borrower with $m=1$ is indifferent between repaying and defaulting on a loan. In this case, a borrower with a good signal is also indifferent between repaying and defaulting, and there is a continuum of equilibria where these two types repay with different probabilities. However, only the equilibrium in which both types, $m=1$ and $m=0$, repay with the same probability, q , is purifiable.¹⁸ We focus on this equilibrium. The probability that a loan made at history B is repaid is now

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

The second equality is established in Section B.1.2 of the Online Appendix. Clearly, $\pi(\tilde{p}, q)$ is a continuous, strictly increasing function of q . Since we are considering the case where $\pi(\tilde{p}, 1) > \frac{\ell}{1+\ell}$, and since $\pi(\tilde{p}, 0) = 0$, there exists a value \tilde{q} setting the repayment probability $\pi(\tilde{p}, \tilde{q})$ equal to $\frac{\ell}{1+\ell}$.

Proposition 4. *Suppose $\bar{K} \geq 2$, and assume the simple information structure with \bar{K} memory. If $0 < \ell < \frac{1}{\bar{K}-1}$, there exists a purifiable mixed equilibrium where the borrower's payoff is strictly greater than \bar{V} . This equilibrium takes the following form. If ℓ is strictly greater than a threshold value ℓ^* , then the borrower plays a pure strategy, where she repays if $m \in \{0, 1\}$. If $\ell \in (0, \ell^*]$, then loans are made with probability \tilde{p} after B , and borrower types $m \in \{0, 1\}$ repay with probability \tilde{q} so as to make the lender indifferent between lending and not lending at B .*

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

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

An Example: We consider parameter values such that $\bar{K} = 4$.¹⁹ Consider the pure strategy profile when $K = \bar{K}$, where the lender extends a loan only after G . The invariant distribution over

18. Section B.1.3 of the Online Appendix proves that the mixed equilibrium where $m \in \{0, 1\}$ repay with the same probability is purifiable. It also shows that there are other, non-purifiable equilibria, that may be more efficient.

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




FIGURE 3

Trust game with costly initiation.

m -values is uniform, and the lender is repaid after lending at B with probability $\frac{1}{4}$. If $\ell > \frac{1}{3}$, a pure strategy equilibrium exists. The expected payoff to a borrower with a good history is $\bar{V} = 0.763$. The lender's expected payoff equals the probability of encountering a borrower with a good history, which is $\bar{W} = 0.714$.

If $\ell < \frac{1}{3}$, lending after B is too profitable and a pure strategy equilibrium with 4-period memory does not exist. A pure strategy equilibrium with longer memory exists, but can be very inefficient. For example, if $\ell = 0.315$, we need $K = 30$, in which case $m^*(K) = 7$. Under the invariant distribution over borrower types, only a quarter of the population has a clean history, and the lender's payoff equals 0.25, strictly less than \bar{W} . Since exclusion is very long, the borrower's payoff at a clean history is also low, at 0.537, which is strictly less than \bar{V} .

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

If ℓ is smaller, say 0.1, then the mixed equilibrium also requires random repayment by borrowers of types $m \in \{0, 1\}$.²⁰ The lending probability after B is $\tilde{p} = 0.034$, and the repayment probability is $\tilde{q} = 0.383$ for $m = 0$ and $m = 1$. The lender's payoff in this equilibrium is substantially lower: $\mu_0(\tilde{q} - \ell(1 - \tilde{q})) = 0.084$. This is largely because a lower fraction of the population has a good history: $\mu_0 = 0.262$. The payoff to the borrower with a clean history equals $V^* = 0.778$.

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

5.3. Initiating trust

We now modify the trust game so that interaction needs to be initiated by the borrower, as in Figure 3. First, the borrower (player 2) can apply for a loan at a small cost, or not. If she does not apply, the game ends, with a payoff $b > 0$ for her. If she applies, the trust game is played. The

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

backwards induction outcome is *out*. Assume the simple information structure with \bar{K} periods memory.²¹

Under this information structure, prior application transforms the interaction between an individual lender and borrower into a signalling game. A borrower with a bad credit history has private information regarding her default incentives which depend on m , the time that must elapse before her credit history becomes good. A borrower who intends to default has a greater incentive to apply. Therefore, in an equilibrium where bad credit risks do not apply, the lender treats any application with suspicion.

Consider the following strategy profile, σ^* . Borrowers with signal G apply, are given a loan, and repay. Borrowers with signal B do not apply; if they do make an application, the lender rejects it; if the lender were to accept their application, the borrower defaults if $m > m^*$ and repays if $m \leq m^*$. The following proposition shows that if $K = \bar{K} \geq 2$, then this strategy profile is a sequentially strict equilibrium, where the lender believes that an applicant with signal B will default with probability one, and these beliefs are implied by the D1 criterion of [Cho and Kreps \(1987\)](#). Thus payoffs are fully efficient, no matter how profitable loans are.

Proposition 5. *Assume that $K = \bar{K} \geq 2$ and the random version of the simple information structure. For any $\ell > 0$, σ^* is a sequentially strict perfect Bayesian equilibrium supported by beliefs assigning probability one to an applicant with a bad credit history defaulting. The borrower beliefs satisfy the D1 criterion. The equilibrium payoffs approximate V^* and W^* .*

6. EXTENSIONS

We have identified the incentives of lenders to lend to defaulting borrowers as a key problem that may undermine the provision of borrower incentives. Of course, governments may directly discourage such lending, by legislation or via central bank regulation. However, such a step may be politically unpopular, especially in the U.S., where there is strong support for a “forgiving” attitude towards bankruptcy and for providing access to credit for borrowers. We now consider some extensions that examine the robustness of our informational mechanism, and how it may be generalized.

6.1. Screening borrowers

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

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

Nevertheless, the deviating lender might offer an alternative contract that uniformly reduces default rates, without inducing separation. For example, the lender might offer a smaller loan at a B signal, that requires a smaller gross repayment, thereby reducing the borrower's incentive to default. Smaller loan sizes induce more borrower types to repay. The deviating lender's profits from making this loan equals the gross repayment times the repayment probability, minus the sum lent and the fixed cost of lending. If the fixed costs of lending are sufficiently large, then reducing loan size may not be profitable, even though this raises the repayment probability. If offering smaller loans is profitable, the punishment for defaults then becomes credit limitation rather than outright exclusion, and the substantive message of our article is unaffected.

Finally, let us suppose that the deviating lender offers a contract with deferred repayment—repayment is due in the next period, rather than at the end of the current period, and has the same present value. This has subtle effects, since borrowers are closer to the end of the punishment phase when they are required to repay. Section B.4 of the Online Appendix shows that deferred repayment may affect the default behaviour of two types of borrower—types $m^* + 1$ and $m = 1$ —but leaves the decisions of other borrower types unchanged. Type $m^* + 1$ may or may not switch her behaviour from defaulting to repaying—the effect is ambiguous. On the other hand, type $m = 1$ will switch to defaulting, since in the next period she gets another loan and has to make two repayments, doubling her incentive to default. Since types are uniformly distributed, conditional on a B signal, offering deferred payments either reduces the overall repayment probability, or leaves it unaffected, so that the deviation is not profitable. These extensions show that our equilibrium is robust to allowing for deviations to alternative contracts.

6.2. Unobserved borrower heterogeneity

We now consider borrower heterogeneity that is not observed by the lenders.²² Suppose that there are two types of borrowers differing only in their rates of involuntary default: normal borrowers with default rate λ and high-risk borrowers with rate $\lambda' > \lambda$. Let r denote the expected payoff of lending to a high-risk borrower who intends to repay, and assume $r > 0$.²³ Let \bar{K}' denote the minimum punishment length required to incentivize repayment by high-risk borrowers. By (1) and (2), the difference $\bar{K}' - \bar{K}$ is positive, increases with the degree of unobserved heterogeneity ($\lambda' - \lambda$), and equals zero if that degree is sufficiently small.

First, we focus on the existence of pure strategy equilibria with minimal punishment length under the simple information structure. If $\bar{K}' = \bar{K}$, unobserved heterogeneity makes it easier to support the efficient pure strategy equilibrium. High-risk borrowers will be over-represented in the pool of borrowers with a bad credit history, reducing the lender's incentive to lend at such a history, and therefore relaxing the crucial incentive constraint.

If unobserved heterogeneity is large enough that $\bar{K}' > \bar{K}$, the information designer faces a trade-off between two sources of inefficiency. Should he increase the punishment length for both types to ensure that the high-risk borrowers have incentives to repay? Or provide efficient punishments for the normal types and incur the social costs of high-risks defaulting? The latter is preferable if the fraction $\eta \in (0, 1)$ of high-risk types in the population is sufficiently small. Then, the pure strategy profile with \bar{K} periods punishment remains an equilibrium, with high-risk types always defaulting, so that the lender's constraint at a bad history is relaxed. For larger η , it may be better to increase the punishment length to \bar{K}' . The complication is that the fraction

22. In what follows, it is not necessary to assume that the borrower knows her type, but for concreteness let us assume that this is the case.

23. From Figure 1, $r := \frac{1 - \lambda' - (\lambda' - \lambda)\ell}{(1 - \lambda)}$, which is less than one, the payoff of lending to a normal borrower.




FIGURE 4
Centipede game.

of normal risks repaying their loan at a bad history might increase, jeopardizing the lender's constraint at a bad history. But that fraction eventually decreases when K becomes sufficiently large, and Proposition 2 ensures that the lender's incentive constraint can be satisfied with longer punishments.

Second, we comment on the relation between unobserved heterogeneity and the unravelling of punishments under *perfect* bounded information on borrower records. If $r > 0$, unravelling still occurs in a candidate pure strategy equilibrium where defaulters are temporarily excluded. But if unobserved heterogeneity is so large that $r \ll 0$, a lender may find it unprofitable to lend to the "average" defaulter ending her punishment, so that unravelling need not occur. In this case, the opposite problem might arise: lenders will find it unprofitable to resume lending to borrowers whose punishments are over, because the absence of loans in a borrower's recent record allows lenders to infer that she has defaulted earlier, and is therefore more likely to be high-risk than a borrower with recent loans. In this case also, there may be social benefits to restricting the information available to lenders.²⁴ Finally, one expects that with improvements in information technology, such as the advent of big data, unobserved heterogeneity would be reduced: lenders will be better able to predict involuntary default risk from observable borrower characteristics. Thus, unravelling is likely to become increasingly important.

6.3. Generalizing our results

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

The class of stage-games that our results extend to includes, in particular, all games where each player moves at most once along any path of play, as only the second mover can have an incentive to deviate. To illustrate, consider the prisoner's dilemma with sequential moves, where the efficient outcome is mutual cooperation. Player 1 cannot profitably defect, since player 2 can respond by defecting. So only player 2 has an incentive to deviate in the one-shot game. But we can also accommodate some games where the players move more than once. In the centipede game in Figure 4, if the players play "across" at each node, this results in payoffs (3, 3). If both

24. This partially echoes the work of Elul and Gottardi (2015) and Kovbasyuk and Spagnolo (2016) discussed in the introduction.

players expect this outcome, then only player 2 has an incentive to deviate. Since moral hazard is one-sided, the analysis of this article applies. However, if player 1's payoff from playing d_3 is changed from 2 to 4, then he would also have an incentive to deviate at this node, and our analysis would not apply.

7. CONCLUSION

We have analysed the trust game played in a large society and examined how moral hazard on the part of borrowers can be deterred. Although moral hazard is one-sided, incentivizing lenders to desist from lending to defaulting borrowers turns out to be non-trivial. Our substantive results show that by providing coarse information, endogenous adverse selection can be leveraged by an information designer to provide such incentives. The information structures and equilibria that sustain efficient outcomes are robust, simple and intuitive.

Our article suggests new directions for empirical work on consumer credit. What is the trajectory of credit limits and loan terms for borrowers prior to bankruptcy flag removal, and how do these compare with terms for borrowers whose flags have just been removed? How likely are such borrowers to default in future, compared to similar borrowers without past defaults, and compared to borrowers whose flags have just been removed? Finally, what are the social welfare consequences of existing credit scoring systems? The stated purpose of credit rating agencies is to help individual lenders predict delinquency—not to sustain socially efficient outcomes. These two goals may well be in conflict.

A. APPENDIX

A.1. Proofs related to Section 3.1

Let V^* be the supremum value of the borrower in any equilibrium. If the borrower gets a loan and repays when her value is near V^* , this must satisfy the inequality:

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

The right-hand side is derived as follows. When the borrower is able to repay, she gets no more than V^* tomorrow, this being the supremum value. V^P denotes the borrower's value following involuntary default.

In any period when her value is near V^* , the borrower must receive a loan, and the lender will only lend if the borrower repays with positive probability. Since the borrower's mixed strategies are not observable, repaying for sure must be optimal, implying the incentive constraint:

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

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

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

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

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

To see that the above upper bound is achievable, let the difference $V^* - V^P$ be such that A.2 holds with equality. This implies that A.1 and therefore A.3 hold with equality.

A.2. Proof of Proposition 1

Fix a sequentially strict equilibrium of the game with perfect K period memory for the lender. Since the borrower observes her entire history, her strategy is a sequence $\sigma := (\sigma_t)_{t=1}^\infty$, where $\sigma_t: \mathcal{O}^{t-1} \rightarrow \{R, D\}$. Let $\rho: \mathcal{O}^K \rightarrow \{Y, N\}$ denote the pure strategy of an arbitrary lender.²⁵

25. We can augment \mathcal{O} with a “dummy” outcome to describe histories in the initial K periods of the game. We do not assume that all lenders/borrowers follow the same strategy, but to avoid needless notation, we do not index these strategies by individual identities.

For an arbitrary natural number T and any integer j , $0 \leq j \leq T$, define the equivalence relation \sim_j on \mathcal{O}^T such that $h \sim_j h'$ iff the last j outcomes in h and h' are identical.

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

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

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

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

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

Since defaulting is optimal at h' ,

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

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

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

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

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

Now let us consider the perturbed game. The payoff to the lender from lending is augmented by εy , where y is the realization of a random variable that is distributed on a bounded support, say $[0, 1]$ (without loss of generality) with a continuous cumulative distribution function, F_Y . The expected payoff to the borrower from wilful default is augmented by εz , where z is the realization of a random variable that is distributed on $[0, 1]$ with a continuous cumulative distribution function, F_Z .

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

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

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

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

where $\tilde{V}_{t+1}(h, \omega)$ denotes the borrower's ex-ante value function at history (h, ω) , before her payoff shock in period $t+1$ is realized. However, since the left-hand side of (A.4) is strictly increasing in z , (A.4) can hold for at most one value of z in $[0, 1]$. This establishes step 2 in the argument.

Turning to step 3, assume that the strategy of the borrower is measurable with respect to \sim_j , with $j \leq K$, and let $h \sim_j h'$. Let $z^*(h)$ denote the value of z such that the borrower is indifferent between repaying and defaulting at h , with $z^*(h) = 0$

(resp. 1) if the borrower always prefers to default (resp. repay). Thus the payoff to the lender from lending at h is

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

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

A.3. Proof of Proposition 2

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

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

Since

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

which proves the required result.

A.4. Proofs Related to Section 5.1

Let $K = \bar{K} \geq 2$. We prove that default is strictly optimal for any $m > 1$. (This implies, in particular, that $\tilde{K} = \bar{K}$ when $\bar{K} > 1$.) That is, we show that for every $m > 1$,

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

Since $V^{\bar{K}}(m)$ is strictly decreasing in m , it suffices to prove this for $m=2$:

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

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

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

Thus, to prove (A.5), it suffices to show that

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

Since $V^{\bar{K}-1}(0) > V^{\bar{K}}(0)$, the inequality (A.5) is established.

A.5. Proof of Proposition 3

Suppose that under the random information structure, the borrower is excluded with probability x for \bar{K} periods and with probability $(1-x)$ for $(\bar{K}-1)$ periods. The borrower's value function at a clean history has the same form as before:

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

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

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

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

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

The borrower agrees to repay at G if and only if

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

is positive.

Using the expression in (A.6), the derivative of the right-hand side of (A.7) with respect to x is

$$\frac{(1-\delta)^3(1-\lambda)\delta^{\bar{K}}}{\left(1-\delta(1-\lambda)-\lambda\delta^{\bar{K}}(x\delta+1-x)\right)^2} > 0, \quad x \in [0, 1].$$

Thus (A.7) is a strictly increasing function of x . By the definition of \bar{K} , it is positive when $x=1$, and strictly negative when $x=0$, and so there exists $x^*(\bar{K}) \in (0, 1]$ setting (A.7) equal to zero. Thus repayment is strictly optimal at signal G when $x \in (x^*(\bar{K}), 1]$.²⁶

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

For every $x \in [x^*(\bar{K}), 1]$, $m^\dagger(\bar{K}, x)$ is the unique value of $m \in (0, K)$ setting

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

equal to zero. The above is a strictly decreasing function of each of the arguments, m and x —the latter follows from the fact that $V^{\bar{K},x}(0)$ is positive and is also strictly decreasing in x . Thus $m^\dagger(\bar{K}, x)$ is a strictly decreasing function of x .

We have already established in Appendix A.3 that $m^\dagger(\bar{K}, 1) \in (1, 2)$. Now we argue that $m^\dagger(\bar{K}, x^*) = 1$, which would prove that $m^\dagger(\bar{K}, x) \in (1, 2)$ for $x \in (x^*, 1]$, since $m^\dagger(\bar{K}, x)$ is strictly decreasing in x . Observe that from the definition of x^* , type $m=0$ is indifferent between repaying and defaulting, since 1 holds with equality. Since type $m=1$ has identical incentives to type $m=0$, she is also indifferent, which implies $m^\dagger(\bar{K}, x^*) = 1$.

Acknowledgments. Thanks to Andrew Atkeson, Dirk Bergemann, Dean Corbae, Mehmet Ekmekci, Mark Feldman, John Geanakoplos, Andy Glover, Johannes Hörner, George Mailath, Larry Samuelson, Tom Wiseman, and seminar audiences at Austin, Chicago, Harvard-MIT, NYU, Stanford, Turin, Toronto, Toulouse, Western Ontario, and Yale for helpful comments. We thank four anonymous referees for many useful suggestions. We are grateful to the Cowles Foundation at Yale for its hospitality while this article was written. V.B. thanks the National Science Foundation for its support via grant #SES 1629055.

Supplementary Data

Appendix B is available at *Review of Economic Studies* online.

REFERENCES

- ALI, S. N. and MILLER, D. A. (2013), “Enforcing Cooperation in Networked Societies” (Mimeo).
- BHASKAR, V. (1998), “Informational Constraints and the Overlapping Generations Model: Folk and Anti-Folk Theorems”, *Review of Economic Studies*, **65**, 135–149.
- BHASKAR, V., MAILATH, G. J. and MORRIS, S. (2013), “A Foundation for Markov Equilibria in Sequential Games with Finite Social Memory”, *Review of Economic Studies*, **80**, 925–948.
- BHASKAR, V. and OBARA, I. (2002), “Belief-Based Equilibria in the Repeated Prisoners’ Dilemma with Private Monitoring”, *Journal of Economic Theory*, **102**, 40–69.
- BHASKAR, V. and THOMAS, C. (2018), “Community Enforcement of Trust” (Mimeo, University of Texas at Austin). Available at <https://sites.google.com/site/carotho14/research/stable/bounded-memory-29jan18.pdf>.
- BULOW, J. and ROGOFF, K. (1989), “A Constant Recontracting Model of Sovereign Debt”, *Journal of Political Economy*, **97**, 155–178.
- CAMERA, G. and CASARI, M. (2009), “Cooperation Among Strangers Under the Shadow of the Future”, *The American Economic Review*, **99**, 979–1005.
- CHO, I.-K. and KREPS, D. (1987), “Signaling Games and Stable Equilibria”, *Quarterly Journal of Economics*, **102**, 179–221.
- DEB, J. (2008), “Cooperation and Community Responsibility: A Folk Theorem for Repeated Random Matching Games” (Discussion paper, mimeo).
- DEB, J. and GONZÁLEZ-DÍAZ, J. (2010), “Community enforcement beyond the prisoner’s dilemma” (Mimeo).
- DOBBIE, W., GOLDSMITH-PINKHAM, P. and MAHONEY, N. (2016), “Bad Credit, No Problem? Credit and Labor Market Consequences of Bad Credit Reports” (Discussion paper, National Bureau of Economic Research).
- DORASZELSKI, U. and ESCOBAR, J. (2010), “A Theory of Regular Markov Perfect Equilibria in Dynamic Stochastic Games: Genericity, Stability, and Purification”, *Theoretical Economics*, **5**, 369–402.

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

- DUFFY, J. and OCHS, J. (2009), "Cooperative Behavior and the Frequency of Social Interaction", *Games and Economic Behavior*, **66**, 785–812.
- EKMEKCI, M. (2011), "Sustainable Reputations with Rating Systems", *Journal of Economic Theory*, **146**, 479–503.
- ELLISON, G. (1994), "Cooperation in the Prisoner's Dilemma with Anonymous Random Matching", *Review of Economic Studies*, **61**, 567–588.
- ELUL, R. and GOTTARDI, P. (2015), "Bankruptcy: Is It Enough to Forgive or Must We Also Forget?", *American Economic Journal: Microeconomics*, **7**, 294–338.
- ELY, J. C. and VÄLIMÄKI, J. (2002), "A Robust Folk Theorem for the Prisoner's Dilemma", *Journal of Economic Theory*, **102**, 84–105.
- FUDENBERG, D. Kreps, D., and MASKIN, E. (1990), "Repeated Games with Long-Run and Short-Run Players", *Review of Economic Studies*, **57**, 555–574.
- GHOSH, P. and RAY, D. (1996), "Cooperation in Community Interaction Without Information Flows", *The Review of Economic Studies*, **63**, 491–519.
- GOVINDAN, S., RENY, P. J., and ROBSON, A. J. (2003), "A Short Proof of Harsanyi's Purification Theorem", *Games and Economic Behavior*, **45**, 369–374.
- GREIF, A. (1993), "Contract Enforceability and Economic Institutions in Early Trade: The Maghribi Traders' Coalition", *American Economic Review*, **83**, 525–548.
- GROSS, T., NOTOWIDIGDO, M. J., and WANG, J. (2016), "The Marginal Propensity to Consume Over the Business Cycle" (Discussion paper, National Bureau of Economic Research).
- HARSANYI, J. C. (1973), "Games with Randomly Disturbed Payoffs: A New Rationale for Mixed-Strategy Equilibrium Points", *International Journal of Game Theory*, **2**, 1–23.
- HELLER, Y. and MOHLIN, E. (2017), "Observations on cooperation", *Review of Economic Studies*, forthcoming.
- KAMENICA, E. and GENTZKOW, M. (2011), "Bayesian persuasion", *The American Economic Review*, **101**, 2590–2615.
- KANDORI, M. (1992), "Social Norms and Community Enforcement", *The Review of Economic Studies*, **59**, 63–80.
- KARLAN, D., MOBIUS, M., ROSENBLAT, T. ET AL. (2009), "Trust and Social Collateral", *The Quarterly Journal of Economics*, **124**, 1307–1361.
- KLEIN, B. and LEFFLER, K. B. (1981), "The Role of Market Forces in Assuring Contractual Performance", *Journal of Political Economy*, **89**, 615–641.
- KOVBASUYUK, S. and SPAGNOLO, G. (2016), "Memory and Markets", (Mimeo) Available at SSRN 2756540.
- LIU, Q. and SKRZYPACZ, A. (2014), "Limited Records and Reputation Bubbles", *Journal of Economic Theory*, **151**, 2–29.
- MUSTO, D. K. (2004), "What Happens When Information Leaves a Market? Evidence from Postbankruptcy Consumers", *The Journal of Business*, **77**, 725–748.
- NAVA, F. and PICCIONE, M. (2014), "Efficiency in Repeated Games with Local Interaction and Uncertain Local Monitoring", *Theoretical Economics*, **9**, 279–312.
- PICCIONE, M. (2002), "The Repeated Prisoner's Dilemma with Imperfect Private Monitoring", *Journal of Economic Theory*, **102**, 70–83.
- RAHMAN, D. (2012), "But Who Will Monitor the Monitor?" *American Economic Review*, **102**, 2767–2797.
- SEKIGUCHI, T. (1997), "Efficiency in Repeated Prisoner's Dilemma with Private Monitoring", *Journal of Economic Theory*, **76**, 345–361.
- SPERISEN, B. (2018), "Bounded Memory and Incomplete Information", *Games and Economic Behavior*, forthcoming.
- STIGLITZ, J. E. and WEISS, A. (1981), "Credit rationing in markets with imperfect information", *The American Economic Review*, **71**, 393–410.
- SUGAYA, T. (2013), "Folk Theorem in Repeated Games with Private Monitoring" (Mimeo, Stanford Graduate School of Business).
- TAKAHASHI, S. (2010), "Community Enforcement When Players Observe Partners' Past Play", *Journal of Economic Theory*, **145**, 42–62.
- TIROLE, J. (1996), "A Theory of Collective Reputations (with Applications to the Persistence of Corruption and to Firm Quality)", *The Review of Economic Studies*, **63**, 1–22.
- WOLITZKY, A. (2012), "Cooperation with Network Monitoring", *Review of Economic Studies*, **80**, 395–427.