Optimal Exclusion

Cyril Monnet
University of Bern and SZ Gerzensee

Erwan Quintin*
Wisconsin School of Business

December 4, 2019

Abstract

We study efficient exclusion policies in a canonical credit model that features both exogenous and strategic default along the equilibrium path. Policies that maximize welfare in a stationary equilibrium implement exclusion for a finite and deterministic number of periods following default. Front-loading exclusion makes the mass of socially valuable transactions as high as it can be in steady state. Less intuitively, doing so also maximizes the average welfare of excluded agents in equilibrium conditional on the level of incentives provided by the threat of exclusion. We argue that these results are robust to a host of natural variations on our benchmark model.

Keywords: endogenous default, exclusion

JEL codes: D52; D82

*Email: cyril.monnet@vwi.unibe.ch and equintin@bus.wisc.edu. We wish to thank Dean Corbae, Briana Chang, Piero Gottardi, Michel Habib, Mark Ready, Stephanie Schmitt-Grohe, Martin Szydlowski, as well as seminar and conference participants at the 2018 Midwest Macroeconomic meetings, the Society for the Advancement of Economic Theory meetings, the North American Meetings of the Econometric Society, and the Society for Economic Dynamics meetings for many useful comments. All remaining errors are our own.

1 Introduction

A large literature studies the role of the threat of exclusion from financial markets in models with endogenous default. Somewhat surprisingly, most papers in this literature take the specification of the exclusion policy as given.[1] For example, in their seminal work on endogenous incompleteness, Kehoe and Levine (1993, 2001, 2008) assume for the most part that exclusion from credit markets is permanent, even though they recognize that other policies could raise welfare.[2] Chatterjee et al. (2006) quantify the effects of various bankruptcy designs in a model where exclusion ends with a positive probability every period. Tertilt et al. (2007) and Liu and Skrzypacz (2013) assume that agents are excluded for a deterministic number of periods. Elul and Gottardi (2015) show that partial exclusion – whereby defaulting agents are only excluded with a certain probability – is generally welfare improving in a model with moral hazard and endogenous borrower reputation.[3]

Our goal in this paper is to characterize the optimal shape of exclusion policies in a canonical model of credit with endogenous and exogenous default. Randomly matched investors and lenders play an ultimatum game over financing terms for a risky investment. Financing contracts feature payment promises from investors to lenders. Investors may fail to pay lenders either because the project fails or by choice so that our model features both strategic and non-strategic defaults in equilibrium. As in Kocherlakota and Wallace (1998) or Bethune et al. (2018), public records make it possible to preserve some information about past payments which allows for dynamic consequences of default and potentially increases the set

[1]Three exceptions are Bond and Krishnamurthy (2004), Corbae et al. (2016) and Kirpalani (2018). Bond and Krishnamurthy (2004) model exclusion as direct constraints on transfers from lenders to investors and study the optimal shape of those constraints. In a model like ours without savings or endowment, their policy would imply permanent exclusion. Corbae et al. (2016) propose a model where exclusion from credit markets arise endogenously as borrowers rebuild their reputation following default. They use their model to quantify the value of having a good reputation in competitive credit markets. In Kirpalani (2018) intermediaries choose to ban agents from financial markets following default instead of simply stopping trade with them because the former approach, unlike the latter, prevents agents from contracting with other intermediaries, hence gives strong truthful revelation incentives.

[2]In section 7, they show by example that exclusion lotteries can raise welfare. Alvarez and Jerman (2000) implement the resulting equilibrium allocations in an environment with portfolio constraints and study the asset pricing properties of the resulting model.

[3]Gu et al. (2013a,b) use probabilistic exclusion and show that it can generate exotic dynamics. Bethune et al. (2018) argue that partial exclusion may be optimal due to a pecuniary externality in an environment where more trade in a particular period tightens borrowing constraints in earlier periods. Zhu (2013) also finds that finite exclusion may be optimal in the dynamic moral hazard model of DeMarzo and Sannikov (2006). In the context of supporting non-cooperative collusion with private information, Green and Porter (1984) show that short periods of non-cooperative plays are necessary to maintain cooperation.

of equilibrium that can be sustained. We consider binary records – investors are either in good or in bad standing – and show that equilibrium may feature partial or full exclusion from borrowing following default until investors' records are cleared. Our key result is that, from a stationary welfare point of view, efficiently designed public records must implement exclusion for a finite but deterministic number of periods. This "harsh but short" punishment approach has two distinct advantages. First, front-loading punishment makes the mass of active investors hence the mass of socially valuable transactions as high as it can be. Second and much less intuitively, this policy maximizes the average welfare of excluded agents in stationary equilibrium. Maximizing the welfare of the excluded also maximizes the number of mutually beneficial transactions in stationary equilibrium so that front-loading punishment is unambiguously efficient.

Front-loaded exclusion policies are a reasonable approximation of how default is punished in practice. In most industrialized nations, one of the primary consequences of credit default by individuals, firms, and sovereigns alike is the *temporary* exclusion from credit markets. On the domestic side, most countries have regulations that allow credit bureaus to record failure-to-pay events and sell that information to creditors. Empirical research has shown that the ability of consumers to borrow is severely impaired by bad records. Bad records, that is, do lead to the effective exclusion of potential borrowers from credit markets. Exclusion, however, is temporary both in practice and as a result of legal constraints for all defaulting agents, be they individuals, firms, or sovereigns. As documented for instance by Elul and Gottardi (2015), most nations impose a statute of limitation that caps the length of credit records. Sovereigns, likewise, experience exclusion following default but are typically able to return to credit markets after a few years.

In the extant literature, making punishment for default as harsh as feasible often raises welfare. Kehoe and Levine (1993, 2001, 2008) describe a dynamic general equilibrium model where the threat of exclusion from credit markets is necessary to support lending along the equilibrium path. In their model, the harsher the exclusion policy, or the higher the consequences of exclusion, the more contracts can typically be supported in equilibrium.[4] In a similar vein, Kocherlakota (1996) considers a dynamic risk-sharing game between two agents with risky endowments. He shows that a feasible allocation in his model can be supported as a subgame perfect equilibrium if and only if at every history each agent receives at least

[4] Krueger and Perry (2005) use this key property to argue that increases in income inequality can lead to less consumption inequality since the penalty associated with exclusion is higher in environments with high income uncertainty.

the utility she would expect in autarky. One interpretation of this result is that optimal subgame perfect arrangements are supported by the threat of permanent exclusion. In his environment, this maximal threat makes the set of sustainable contracts as large as it can be, hence is optimal. Our model does not have this property: extreme punishments are usually suboptimal. Like borrowers do in practice, our investors default in some cases because they have no choice, while others choose to default even though they could pay what they owe. This maps neatly into what applied economists typically classify as strategic and non-strategic defaults.[5] Default rates, therefore, are bounded below. This means that exclusion has to be finite in length almost surely for there to be positive trade in any stationary equilibrium. Even when infinite exclusion is feasible, it is never efficient. In fact, as we mentioned above, efficient exclusion policies minimize the duration of exclusion – instead of maximizing it – conditional on the level of punishment needed to deter strategic default.[6]

Our result resembles standard findings in the classical analysis of repeated game with discounting. Abreu (1988) shows that discounted games can be completely analyzed using "simple" strategy profiles which specify a path of preferred actions and punishments for any deviation from that path. Like in our model, the shape of punishment is independent of history and, optimally, decreases in harshness through time. In that setting, lowering the severity of punishment over time is necessary for subgame perfection: "early stages of an optimal punishment must be more unpleasant than the remainder [...] to deter a player from cheating when he is already being punished as harshly as possible." In our case, it is efficient to front-load punishment because any early forgiveness must be compensated for disproportionately in the future.

While the efficient shape of exclusion policies can be fully characterized, we show that the optimal length of exclusion depends in ambiguous ways on model parameters. For instance, we present versions of our model in which more project risk is associated with longer exclusion while the exact opposite happens in other, equally reasonable versions. Put another way, while

[5]See e.g. Foote et al. (2008) and Gerardi et al. (2015).

[6]Elul and Gottardi (2015) find that forgiveness can be optimal in a model of borrower reputation. In our model, eventual forgiveness with probability one must be optimal, for otherwise all agents would asymptotically find themselves excluded. This aspect of our model is similar in spirit to a point made by Dubey et. al (2005) and Quintin (2013) about the optimal intensity of direct default punishment. In both papers, the set of contracts borrowers and lenders can write is exogenously restricted and, as a result, default is a part of equilibrium outcomes. In those environments, punishing default more harshly can lower welfare. It can even lead to higher default rates and, like in our environment, typically leads to fewer transactions. Our paper focuses on exclusion threats and forgiveness rather than direct punishment, but it does share the feature that maximizing the punishment via exclusion would lead to eliminating all lending.

the efficient shape of exclusion can be characterized with remarkable generality, the efficient length of exclusion is highly sensitive to modeling details.

2 The environment without public records

Consider an economy in which time is discrete and infinite. There is one good that cannot be stored across periods. A mass one of infinitely-lived investors are each endowed with a project but no good. They can activate their project in each period by investing a quantity $k \in K$ of the good at the start of a period where K is the set of feasible scales of operation. For instance, it may be that $K = [\underline{k}, +\infty]$ in which case $\underline{k} \geq 0$ is a minimal scale of operation. Or K may be a finite set, even a singleton, as will be the case in some of the examples we consider in the sequel.

When it is successful, which occurs with probability $\pi > 0$ in each period, the project pays output Ak at the end of the period where $A > 0$. With the complementary probability, the project pays nothing. A law of large numbers holds so that π is also the fraction of projects that deliver positive output in a given period.

To simplify the exposition, we will first assume that investors are risk-neutral and discount future payoffs at a constant period rate of $\beta \in (0, 1)$. We relax this linearity assumption in section 5. There we consider an environment with risk-averse investors and traditional risk-sharing contracts. While this makes notation and derivations more burdensome, our main results are unchanged.

The economy also contains a mass one of lenders endowed with an amount $M > 0$ of the unique good. Lenders consume at the end of the period, have linear preferences, and discount future flows at the same rate β as investors. They can store the endowment they receive at the start of the period for a time-invariant and safe payoff $R \in (0, A)$ which exogenously pins down the opportunity cost of their resources. In addition, each lender is randomly matched each period with exactly one investor. In any match at date t, lenders have the option to lend a quantity $k_t \in [0, M] \cap K$ of the good to their counterpart, while investors have the option to make a transfer $m_t \in [0, Ak_t]$ to the lender when their project succeeds. Investors make a take-it-or-leave-it offer to the lender with whom they are matched. We will consider different offer protocols in section 5.1. In this section, we assume that lenders have no information about past actions by investors. Our upcoming results will consider the role of providing such information.

Only investors observe their output realization. Furthermore, commitment on their part is limited. If they abstain from making a payment when one is due they experience a penalty $\tau \geq 0$ drawn at the end of the period from a known distribution F. The default cost τ captures in simple fashion any remorse or shame defaulting investors experience as well as the effects of exogenous punishments in the spirit of Dubey et al. (2005). One can also think of τ as the cost associated with hiding or absconding with the project's output. In section 5.2 we will consider the case where default entails not just private costs but also social costs and the case where lenders also experience a private cost when default occurs.

To simplify the analysis we assume that F is absolutely continuous over a bounded support $[0, \bar{\tau}]$. The fact that ex-post default costs can be non-degenerate allows for equilibrium in which a strictly interior fraction of agents default for strategic reasons but plays no other role in our analysis. In fact, we could abstract from direct default costs altogether without any impact on our main results.[7]

Absent any dynamic consequences of default, investors of ex-post type $\tau \geq 0$ repay their loan when and only when their project is successful and

$$m_t \leq \tau.$$

Note that payments cannot be made contingent on τ because we assume that it is drawn after offers are made. In section 5.3 we study the case with adverse selection where τ is drawn before offers are made. Payments are contingent on output however but to ease notation we do not make that dependence explicit since payments are clearly zero when the project fails so that any positive payment stipulation is de facto contingent on success.

Our environment captures in a simple way the two types of default events that correspond to the classification used in applied work.[8] Strategic default occurs when the project is successful but, nonetheless, investors choose to experience disutility τ rather than honor their debt. Non-strategic, exogenous default occurs when the project fails and investors have no choice but to default. Only investors know why they failed to pay so that other agents cannot distinguish defaults types from one another.

From the point of view of lenders, the expected payoff on loans to investors exceed storage

[7] Allowing F to feature jumps (as we do in various examples in section 6) would not change any of our results but equilibrium may then feature mixed default strategies for positive masses of investors.

[8] See e.g. Foote et al. (2008).

returns when:

$$\pi(1 - F(m_t))m_t \geq k_t R. \tag{2.1}$$

Indeed, only the mass $1 - F(m_t)$ of successful investors whose τ is below m_t repay. This is the participation constraint for lenders in any match. As Quintin (2012, 2103) discuss at length, $(1 - F(m_t))m_t$ is generally non-monotonic. Furthermore, because F has bounded support, the right-hand side of the last constraint is bounded by $\pi(1 - F(m^*))m^*$ where m^* is a maximizer of $(1 - F(m))m$. This implies that feasible capital advances are bounded in any match. Investors select m_t and k_t to solve

$$(P1) \qquad \max_{k_t, m_t} \pi \left(Ak_t - \int_0^{m_t} \tau dF + \int_{m_t}^{\bar{\tau}} m_t dF \right)$$

subject to:

$$\begin{aligned} k_t &\in K, \\ k_t &\leq M, \\ m_t &\leq Ak_t, \\ k_t R &\leq \pi(1 - F(m_t))m_t. \end{aligned}$$

The feasible set is empty when no capital level compatible with payment feasibility and the lenders' participation constraint is in K. In all other cases a solution to this problem always exists since the feasible set is compact while the investors' objective function is continuous. While several solutions may exist, the final constraint (condition 2.1, that is) must bind at any solution. Indeed, were it not the case, it would be feasible to either raise k_t or lower m_t until the constraint does bind which would strictly raise the objective. It follows that solutions are fully summarized by m_t since the associated capital use is implied by condition 2.1. Furthermore, whenever a project is activated, we must have $m_t > 0$ since condition 2.1 could not hold otherwise for any $k_t > 0$. So we will write $m_t = 0$ when no contract is feasible.

Since we have yet to introduce a public record technology there are no meaningful dynamics in this version of our economy. It is essentially a special case of the two-date model Dubey et al. (2005) describe in which all default punishment is a direct and exogenous utility penalty. Not surprisingly then, existence holds in our model in as much generality as it does in theirs. Specifically, every period, investors propose to lenders a solution to problem $(P1)$ when one exists. When no solution exists, lenders simply store their endowments each period while investors do not activate their project.

3 Public records and exclusion

Assume now that partial information about past actions by investors is available in the form of an investor-specific public signal $s \in \{B, G\}$. For reasons that will soon become transparent, we will think of investors whose status at the start of a period is $s = G$ as investors in good standing while investors whose $s = B$ are in bad standing. Consider first investors who begin a period in good standing. If they enter into a loan contract with the lender with whom they are matched that features $m_t > 0$ and fail to pay either by choice or because the project failed, their status becomes $s = B$ at the end of the period. Otherwise, they remain in good standing. If, on the other hand, investors enter a period in bad standing, the evolution of their public status is governed by a sequence $\{\phi_s\}_{s=0}^{+\infty}$ of *forgiveness* probabilities. With probability ϕ_0 his status is immediately – before they become matched – adjusted to $s = G$. After this first draw, investors who have been in bad standing for n periods are forgiven with probability ϕ_n.

In other words, for all n, ϕ_n is the hazard rate into return to good standing. The probability that an investor is going to be in bad standing for exactly n periods following default is

$$\phi_{n+1} \prod_{s=0}^{n} (1 - \phi_s).$$

By the same logic,

$$\prod_{s=0}^{+\infty} (1 - \phi_s)$$

is the probability that investors in bad standing will never be forgiven.

This general formulation encompasses all the specifications employed in the existing literature on endogenous default. Immediate forgiveness ($\phi_0 = 1$) yields the economy we described in the previous section, a special case of the environment described by Dubey et al. (2005). At the opposite extreme, zero forgiveness ($\phi_s = 0$ for all s) – which results in permanent exclusion – is the case studied in Kehoe and Levine (1993). Chatterjee et al. (2006) specify constant forgiveness rates, i.e. $\phi_s = \phi$ for all s. The one-time forgiveness lottery of Elul and Gottardi (2014) or section 7 in Kehoe and Levine (1993) corresponds to the case where $\phi_0 > 0$ but $\phi_s = 0$ for all $s > 0$. Setting $\phi_s \in \{0, 1\}$ for all s gives exclusion for a deterministic number of periods as in Tertilt et al. (2007) or Liu and Skrzypacz (2013). Our formulation also allows in principle for much more complex forgiveness policies.[9]

Investors can now be in different states at the start of a given period. Some investors are in good standing. Among investors who are not in good standing, we also need to record the number of periods for which they have been in bad standing since, in general, the forgiveness probability may depend on that length of time. Investor status and the number of periods for which investors have been in bad standing are public information. This makes our environment similar to the incomplete record-keeping model of Kocherlakota and Wallace (1998) and, like them, we consider symmetric and stationary equilibrium (SSE's) given a forgiveness policy. In each match, investors and lenders play an ultimatum game in which investors propose a contract (k, m) and lenders accept or reject it. Rejections lead to zero payoff for investors while lenders store their endowment and earn MR.

Equilibria specify actions for investors and lenders in any match for every possible action of their counterpart, not simply actions that are played along the equilibrium path. In particular, lenders announce a complete acceptance policy. Actions may only depend on the investors' status of record, namely whether or not they are in good standing and, for investors who are in bad standing, how many periods they have been in that state. The stationarity requirement is that actions may not depend on the date,[10] while the symmetry requirement is that counterparts must treat all possible matches equally, so that actions cannot depend on an investor's identity.

It is well known that an environment like ours may produce many types of equilibria, even after restricting attention to stationary and symmetric strategies. We will consider Markov perfect equilibria in which investors in good standing offer lenders a contract (k_G, m_G) while investors in bad standing offer lenders a potentially different contract (k_B, m_B). Towards constructing such equilibrium, observe first that with any quadruple (k_B, m_B, k_G, m_G) of actions are associated continuation values V^G and $V^B(n)$ for investors in each possible state. Taking investors in good standing first,

[9] We are making the implicit assumption that investors are all treated equally once they default. This is without loss of generality as long as the default shock is independent across periods. But if types are persistent, welfare may be improved by making punishment depend on the credit history of investors. See section 5.3.

[10] Our type of environments could support time-dependent strategies and time-dependent forgiveness policies. One could imagine for instance periods of limited lending and limited punishment alternating with periods of high lending and high punishment, which would generate lending cycles in the sense of Gu et al. (2013b) or Bethune, Hu, and Rocheteau (2018). Focusing on SSEs is a limitation in the sense that it rules out such cycles.

$$V^G = (1-\pi)\beta V^B(0) + \pi E_\tau \max\left\{Ak_G - m_G + \beta V^G, Ak_G - \tau + \beta V^B(0)\right\} \quad (3.1)$$

while for investors who have been in bad standing for $n \geq 0$ periods,

$$V^B(n) = \phi_n V^G + (1-\phi_n)\left[\pi E \max\{Ak_B - m_B, Ak_B - \tau\} + \beta V^B(n+1)\right]. \quad (3.2)$$

To understand this second expression, recall that investors in bad standing improve their standing if and only if they are forgiven. While they may receive a loan while in bad standing, whether or not they repay this loan has no bearing on their standing. In section 5.7, we will briefly discuss the case where investors in bad standing may improve their standing by entering into a loan and repaying it. Importantly,

Lemma 3.1. *Continuation values V^G and $V^B(n)$ are uniquely defined for all n for any set of contract proposals (k_B, m_B, k_G, m_G) and given a forgiveness policy $\{\phi_s\}_{s=0}^{+\infty}$.*

Proof. Given (k_B, m_B, k_G, m_G), a guess for $V^B(0)$ implies a unique value for V^G by equation 3.1. In turn, equation 3.2 implies a unique value of $V^B(0)$. To see this, relying on a traditional value function iteration approach, posit $V^B(N) = 0$ for a large N. Backwards induction then gives a unique approximation for $V^B(n)$ for $n \leq N$, including for $V^B(0)$. Because $\beta \in (0,1)$ this approximation for $V^B(0)$ converges to a unique value as N gets large. This yields a mapping on the real line which takes any guess for $V^B(0)$ into a new guess. Standard arguments show that this mapping defines an operator on the real line that satisfies Blackwell's sufficient conditions for a contraction. In particular, it has a unique fixed point. Since $V^B(0)$ is uniquely defined, so are V^G and, in turn, $V^B(n)$ for all n. $\square$

Turning now to individual rationality, investors choose to pay if the project succeeds and once they discover their default cost τ if

$$-m_G + \beta V^G \geq -\tau + \beta V^B(0) \iff m_G \leq \tau + \beta\left[V^G - V^B(0)\right] \quad (3.3)$$

so that it is incentive feasible for lenders to accept a proposal (k_G, m_G) in a match if and only if

$$\pi\left(1 - F\left(m_G - \beta\left[V^G - V^B(0)\right]\right)\right) m_G \geq k_G R. \quad (3.4)$$

Importantly, since no information is recorded on lenders' actions in a match outside of the public record, future offers cannot be made conditional on lenders' current actions, so that incentive feasibility for lenders does not involve any dynamic consequences of their acceptance actions. Investors in bad standing choose to repay ex-post if and only if $m_B < \tau$ so that it is individually rational for lenders to accept an offer from bad investors if and only if condition 2.1 holds, i.e.:

$$k_B R \leq \pi(1 - F(m_B))m_B. \tag{3.5}$$

In addition, since default has no dynamic implications for investors in bad standing, their offers solve $(P1)$ in any SSE. More formally, a SSE is a quadruple (k_B, m_B, k_G, m_G), the associated policy functions V^G and $V^B(n)$, and stationary acceptance policies for lenders for every possible match and every possible offer, such that:

1. The acceptance policy is optimal for lenders for every possible offer in any match. That is, offers are accepted if and only if they satisfy (3.3) when made by investors in good standing and (3.4) when made by investors in bad standing;

2. Proposal (k_B, m_B) solves problem $(P1)$;

3. Proposal (k_G, m_G) maximize

$$\pi\left(Ak - \int_0^{m-\beta\left[V^G - V^B(0)\right]} \tau dF + \int_{m-\beta[V^G - V^B(0)]}^{\bar{\tau}} m dF\right) + (1-\pi)\beta V^B(0)$$

   subject to:

$$\begin{aligned} k &\in K, \\ k &\leq M, \\ m &\leq Ak, \\ \pi\left(1 - F\left(m - \beta\left[V^G - V^B(0)\right]\right)\right) m &\geq kR. \end{aligned}$$

All equilibrium have to be such that (k_B, m_B) solves $(P1)$ along the equilibrium path. Having found one such solution, setting $(k_B, m_B) = (k_G, m_G)$ gives a first SSE. Indeed, in that case, $V^B(n) = V^G$ for all n and the participation constraint is exactly the same for investors in good standing and investors in bad standing. In general however, there are many more equilibrium, as we will soon establish.

Any SSE with $m^G > 0$ features a fraction δ^D of investors in good standing who fail to pay where

$$\delta^D = (1-\pi) + \pi F\left(m_G - \beta\left[V^G - V^B(0)\right]\right) > 0.$$

Together with the exogenous forgiveness process, this defines a Markov chain over investor standing. If a stationary equilibrium exists with a constant mass of investors in each standing, the mass μ^G of investors who enter a period in good standing (**before** forgiveness draws are made) must satisfy:

$$\mu^G = \left[\mu^G + \sum_{n=0}^{+\infty} \mu^B(n)\phi_n\right](1-\delta^D) \tag{3.6}$$

where $\mu^B(n)$ is the mass of investors who enter a particular period in bad standing and have been in bad standing for n periods, for all $n \geq 0$. In words, to enter a period in good standing involves being in good standing in the previous period, upon entry or through forgiveness, and not defaulting. As for the mass of investors in bad standing at the start of the period

$$\mu^B(0) = \left[\mu^G + \sum_{n=0}^{+\infty} \mu^B(n)\phi_n\right]\delta^D$$

while, for $n > 0$,

$$\mu^B(n) = \mu^B(n-1)(1-\phi_{n-1}).$$

The following result tells us which forgiveness policies can support stationary equilibrium with a positive mass of investors in good standing.

Proposition 3.2. *A stationary equilibrium with $\mu^G > 0$ exists if only if*

$$\sum_{n=0}^{+\infty} \Pi_{s=0}^{n}(1-\phi_n) < +\infty. \tag{3.7}$$

Proof. Assume by way of contradiction that $\sum_{n=0}^{+\infty} \Pi_{s=0}^{n}(1-\phi_n) = +\infty$ but that a SSE exists with $\mu^G > 0$. For all $n \geq 1$,

$$\mu^B(n) \geq \mu^B(0)\Pi_{s=0}^{n-1}(1-\phi_n).$$

But

$$\mu^B(0) \geq \mu^G\delta^D$$

since the mass of newly excluded investors includes those who where not excluded at the start of the previous period and defaulted. (As we mentioned above, $\mu^B(0)$ also includes borrowers who were forgiven in the previous period and defaulted immediately, but we do not need to count them precisely in the context of this proof.) It follows from these inequalities that:

$$\sum_{n=0}^{+\infty} \mu^B(n) \geq \mu^G \delta^D \sum_{n=0}^{+\infty} \Pi_{s=0}^{n-1}(1-\phi_n) = +\infty.$$

But at any stationary distribution we need

$$\mu^G + \sum_{n=0}^{+\infty} \mu^B(n) = 1.$$

This contradiction establishes necessity. As for sufficiency, we have already argued that setting $(k_G, m_G) = (k_B, m_B)$ where (k_B, m_B) solves problem $(P1)$ is a SSE. If $(k_B, m_B) = (0, 0)$ then no lending ever takes place along the equilibrium path at that SSE and $\mu^G = 1$ is invariant as needed. If $(k_B, m_B) > (0, 0)$ then as long as condition (3.7) holds the countable state Markov Chain is positive recurrent (see Tierney, 1994), hence has a unique invariant distribution, and that distribution must feature $\mu^G > 0$. $\square$

To understand this result, assume for instance that

$$\Pi_{n=0}^{+\infty}(1-\phi_n) \in (0, 1),$$

among other ways in which we may have

$$\sum_{n=0}^{+\infty} \Pi_{s=0}^{n}(1-\phi_n) = +\infty.$$

Then a strictly positive fraction of investors who are in bad standing never return to good standing. This implies that from any initial distribution, the countable state space Markov chain that governs investor standing implies a sequence of distributions such that μ^G converges to zero. In particular, the Markov chain admits no invariant distribution. Hence, the only forgiveness policies that are compatible with strictly positive masses of agents in good standing in equilibrium are the ones for which investors in bad standing are eventually forgiven with probability one.

In the next section, we will concentrate our attention on SSEs in which agents in bad standing are excluded from loan markets. For such equilibrium to exist, it is necessary that parameters be such that $(k_B, m_B) = (0, 0)$ be the only solution to $(P1)$. In that case, investors in bad standing have to be excluded from borrowing until they are forgiven and $(k_B, m_B, k_G, m_G) = (0, 0, 0, 0)$ is one SSE. It is the only SSE when $\phi_0 = 1$ so that investors are immediately forgiven following default. But whenever $\phi_0 < 1$, SSEs with $k_G > 0$ and, for that matter, $\mu^G > 0$ may exist as we will now show by example.

Example 3.3. *Assume temporarily that $\phi_n = 0$ for all n so that exclusion is permanent with probability one. This violates condition 3.7 which simply means that for now all investors are asymptotically excluded but this is not incompatible with the existence of a SSE. By the end of the construction, the exclusion policy will satisfy condition 3.7.*

Also specify $K = [\underline{k}, +\infty]$ for the feasible capital set where $\underline{k} > 0$ is a minimal scale of operation. Denote by m^* the value that maximizes $\pi (1 - F(m))\, m$ and assume that

$$\pi (1 - F(m^*))\, m^* = \underline{k}R - \epsilon$$

where $\epsilon > 0$. In that case, $(k_B, m_B) = (0, 0)$ is the only solution to $(P1)$, which implies that $V^B(0) = 0$ at any SSE. Assume however that

$$\pi m^* \geq \underline{k}R$$

so that a solution to $(P1)$ would exist if default were not an option. So to sustain an SSE with $(k_B, m_B) > (0, 0)$ we simply need exclusion to be sufficiently dissuasive.

To that end, define $V^G(m^*) \geq 0$ to be the only solution to 3.1 given $V^B(0) = 0$ when $(k_G, m_G) = (\underline{k}, m^*)$. Note that $V^G(m^*) \geq \pi (A\underline{k} - m^*)$ which can be made as large as one wants by making A large. In particular, we may choose A large enough so that

$$\pi \left(1 - F\left(m^* - \beta V^G(m^*)\right)\right) m^* \geq \underline{k}R.$$

But this implies that the unique (k_G, m_G) that satisfies condition 2 of the definition of a SSE is such that $k_G > 0$, as needed.

To complete the construction, we need to relax the assumption that $\phi_n = 0$ for all n so as to have built an SSE in which $\mu^G > 0$ at the invariant distribution. Consider the class of forgiveness technologies where $\phi_n = 0$ for all $n < N$ while $\phi_N = 1$. As N grows large, $V^B(0)$

converges towards zero while $V^G(m^*)$ *grows to a value that exceeds* $\pi\left(A\underline{k}-m^*\right)$. *It follows by continuity that*

$$\pi\left(1-F\left(m^*-\beta\left(V^G(m^*)-V^B(0)\right)\right)\right)m^*\geq\underline{k}R$$

so that, once again, $(k_G,m_G)>(0,0)$ *at any SSE, and this SSE is now such that* $\mu^G>0$.

We conclude this section with a key observation that underlies the notion of efficiency we adopt in this paper and contains the intuition for our main results.

Lemma 3.4. *Assume that* (k_G,m_G,k_B,m_B) *is a SSE given forgiveness policy* $\{\phi_n\}_{n=0}^{+\infty}$. *Then* (k_G,m_G,k_B,m_B) *is also a SSE for any other forgiveness policy* $\{\phi_n'\}_{n=0}^{+\infty}$ *that implies the same* $V^B(0)$.

Proof. Given (k_G,m_G,k_B,m_B), knowing $V^B(0)$ implies V^G by condition 3.1. Since (k_G,m_G,k_B,m_B) is a SSE under the original policy, (k_G,m_G,k_B,m_B) remains an SSE under the new policy. □

Our efficiency arguments rely heavily on this fact. Given an SSE (k_G,m_G,k_B,m_B) under a particular forgiveness policy that implies a particular value of default $V^B(0)$, a different forgiveness policy that delivers the same $V^B(0)$ supports the same SSE. Efficient policies then, at a minimum, should deliver the highest average welfare among forgiveness policies that leave $V^B(0)$ unchanged. This minimal requirement turns out to sharply restrict the shape of efficient forgiveness policies, as we explain in the next section. The simplicity of the argument also means that our findings are robust to many variations of our model of borrowing and lending. Once V^G and $V^B(0)$ are known our arguments apply regardless of the economic model that underlies those value functions. We will discuss several of these variations in section 5.

4 Optimal exclusion in stationary equilibrium

To deal with a simple case first we first restrict our attention on SSEs such that investors are excluded from any lending while in bad standing. This will enable us to refer to the forgiveness policy as an exclusion policy. We showed in the previous section that parameters can in fact be restricted so that investors in bad standing do not invest at all while in bad standing so that they are effectively excluded following default until forgiven. In the next section, we will fully generalize our results to the case where $(k_B, m_B) > (0,0)$.

When $(k_B, m_B) = (0,0)$, we obtain via condition 3.2 and given a forgiveness policy $\{\phi_n\}_{n=0}^{+\infty}$ that

$$V^B(0) = \phi_0 V^G + (1-\phi_0)\phi_1\beta V^G + (1-\phi_0)(1-\phi_1)\phi_2\beta^2 V^G + \ldots < V^G \tag{4.1}$$

since

$$\pi E \max\{Ak_B - m_B, Ak_B - \tau\} = 0$$

in that case.

We call a forgiveness policy $\{\phi_n\}_{n=0}^{+\infty}$ efficient if it supports a SSE with invariant distribution $\{\mu^G, \mu^B\}$ and continuation utility $\{V^G, V^B\}$ such that no other policy $\{\phi'_n\}_{n=0}^{+\infty}$ supports the same SSE with invariant distribution $\{\mu_o^G, \mu_o^B\}$ and continuation utility $\{V_o^G, V_o^B\}$ with

$$\mu^G V^G + \sum_{s=0}^{+\infty} \mu^B(s) V^B(s) < \mu_o^G V_o^G + \sum_{s=0}^{+\infty} \mu_o^B(s) V_o^B(s).$$

As we mentioned above, the lenders' participation constraint must bind in any match so that their welfare is $\frac{MR}{\beta}$ in any SSE. Since their welfare is constant across SSEs and forgiveness policies, it cannot have any bearing on efficiency rankings. In section 5.1 we will study equilibria in which lenders earn positive surplus and the efficiency criterion must be amended. The welfare functional we state takes full account of default costs since they are part of the computation of value functions. In section 5.2 we will consider the case where default entails a social cost as well as a private cost.

We further restrict our attention to policies such that

$$\sum_{n=0}^{+\infty} \Pi_{s=0}^{n}(1-\phi_n) < +\infty$$

so that the stationary objective function is well defined. As we have shown, absent this condition our economy must drift towards a distribution of types with $\mu^G = 0$.

Ours is a minimal efficiency requirement in the sense that holding loan terms the same, an efficient policy maximizes the average welfare of agents in steady state.[11] Minimal a requirement as it may seem, it is enough to greatly restrict the set of forgiveness policies.

[11] Appendix 9.3 studies a different notion of efficiency based on welfare as of date 0 rather than stationary welfare.

Theorem 4.1. *Efficient forgiveness policy* $\{\phi_s\}_{s=0}^{+\infty}$ *must be such that for some* s^*,

1. $\phi_s = 0$ *for all* $s < s^*$;

2. $\phi_{s^*} \in (0, 1]$;

3. $\phi_{s^*+1} = 1$ *and* $\phi_s \in [0, 1]$ *for all* $s > s^* + 1$.

Efficient forgiveness policies bring investors in bad standing back into lending markets as fast as possible given that a certain level of punishment needs to be imposed in equilibrium. This is optimal because doing so has two distinct virtues. First, it maximizes the stationary mass of investors in good standing, which is positive for welfare since $V^B(n) < V^G$ for all n. Second and much less intuitively, it also maximizes the average welfare of investors in bad standing. The second feature stems from the fact that if the forgiveness policy is lenient early in the sense that it forgives some of the investors who defaulted right away, it must reestablish incentives by increasing punishment on investors that have been in bad standing for longer periods of time. Discounting implies that this punishment more than offsets the gains of the investors released from exclusion.

Note that the optimal exclusion policy involves a non-degenerate forgiveness lottery in at most one period. In the natural continuous time limit of our environment, the optimal policy would not require this randomization device. An equivalent interpretation of $\phi_{s^*} \in (0, 1]$ is that investors are only allowed to operate their technology for part of transition period s^*. In particular, no randomization device is necessary to implement the optimal exclusion policy.

Towards proving our main theorem, consider any SSE (k_G, m_G, k_B, m_B) supported by a forgiveness policy that implies a certain $V^B(0)$. By lemma 3.4, for the associated policy to be efficient, it must solve:

$$\max_{\{\phi_s\}_{s=0}^{+\infty}} \quad \mu^G V^G + \sum_{s=0}^{+\infty} \mu^B(s) V^B(s)$$

subject to

$$V^B(0) = \phi_0 V^G + (1 - \phi_0)\phi_1 \beta V^G + (1 - \phi_0)(1 - \phi_1)\phi_2 \beta^2 V^G + \dots \tag{4.2}$$

We will show that any solution to the above problem satisfies the conditions listed in theorem

4.1. Observe that

$$\mu^G V^G + \sum_{s=0}^{+\infty} \mu^B(s) V^B(s) = V^G + \sum_{s=0}^{+\infty} \mu^B(s) \left(V^B(s) - V^G\right)$$

since

$$\mu^G + \sum_{s=0}^{+\infty} \mu^B(s) = 1,$$

so that we can simply maximize

$$\sum_{s=0}^{+\infty} \mu^B(s) \left(V^B(s) - V^G\right)$$

subject to 4.2. Since for all $s > 0$,

$$\mu^B(s) = \mu^B(0) \prod_{i=0}^{s-1} (1 - \phi_i)$$

this objective can be re-written as

$$\mu^B(0) \left\{ \left[V^B(0) - V^G\right] + (1 - \phi_0) \left[V^B(1) - V^G\right] + (1 - \phi_0)(1 - \phi_1) \left[V^B(2) - V^G\right] + \ldots \right\}$$

To proceed we consider first the problem of maximizing the bracketed expression

$$\left[V^B(0) - V^G\right] + (1 - \phi_0) \left[V^B(1) - V^G\right] + (1 - \phi_0)(1 - \phi_1) \left[V^B(2) - V^G\right] + \ldots$$

subject to 4.2. In other words, we consider the problem of maximizing the *average* welfare of excluded investors in stationary equilibrium. Using expression (3.2), the resulting objective is

$$\begin{aligned}
\mathcal{P} &= -(1 - \phi_0) \left[V^G - \beta V^B(1)\right] - (1 - \phi_0)(1 - \phi_1) \left[V^G - \beta V^B(2)\right] - \ldots \\
&= -V^G[(1 - \phi_0) + (1 - \phi_0)(1 - \phi_1) + (1 - \phi_0)(1 - \phi_1)(1 - \phi_2) + \ldots] \\
&+ (1 - \phi_0)\beta V^B(1) + (1 - \phi_0)(1 - \phi_1)\beta V^B(2) + (1 - \phi_0)(1 - \phi_1)(1 - \phi_2)\beta V^B(3) + \ldots
\end{aligned}$$

In the appendix we show that the final part of the expression for $\mathcal{P}$ (the last line in the string of equations above) is constant over the constraint set defined by (4.2), as long as

$$\sum_{n=0}^{+\infty} \Pi_{s=0}^{n}(1-\phi_n) < +\infty.$$

While the detailed argument for why this is true is quite involved, the intuition is simple. The final part of the expression is proportional to the utility investors expect following default, a level which is pinned down by the constraint. Given this fact, it follows that maximizing $\mathcal{P}$ over the constraint set amounts to minimizing

$$\zeta \equiv (1-\phi_0) + (1-\phi_0)(1-\phi_1) + (1-\phi_0)(1-\phi_1)(1-\phi_2) + \dots$$

subject to the constraint that the right level of punishment must be imposed.

We show in the appendix that this is done by making early ϕ's zero until the constraint (4.2) is met. In other words, maximizing the average welfare of excluded investors given that punishment (4.2) must be inflicted on investors who just defaulted is optimally done by fully excluding them until they have suffered precisely the punishment equilibrium requires. Then investors return to the non-excluded fold with probability one. The intuition for this result is as follows. One could even out forgiveness chances across excluded investors of each type. For instance, a candidate policy could feature the unique constant ϕ that meets the punishment constraint. This would give defaulting investors a chance to avoid exclusion altogether, for one potential benefit. But then one needs to increase punishment (reduce forgiveness odds) in the future to preserve $V^B(0)$. Because of time-discounting ($\beta < 1$) the increase in future punishment more than undoes the benefits of reducing the severity of punishment in early periods. This intuition is formalized in the variational argument developed in the appendix.

The bottom line is that maximizing the average welfare of the excluded is the same as minimizing ζ. But ζ admits a convenient interpretation for our purposes: it is monotonically related to the mass of excluded investors in any period. Indeed, recall that $\mu^B(0)$ is the mass of agents that just became excluded at the start of a given period. In turn,

$$\mu^B(1) = (1-\phi_0)\mu^B(0),$$

while

$$\mu^B(2) = (1-\phi_0)(1-\phi_1)\mu^B(0)$$

and so on and so forth. It follows that the total mass of excluded may be written as

$$\sum_{s=0}^{+\infty} \mu^B(s) = \mu^B(0)(1+\zeta).$$

In words, minimizing ζ corresponds to minimizing the mass of investors who are excluded in any period. To proceed, the following expression for $\mu^E(0)$ will be useful.

Lemma 4.2. *In any stationary equilibrium with positive investment,*

$$\mu^B(0) = \mu^G \times \frac{\delta^D}{1-\delta^D}$$

where δ^D is the time-invariant default rate on loans.

Proof. In any stationary equilibrium,

$$\mu^B(0) = \left[\mu^G + \sum_{n=0}^{+\infty} \mu^B(n)\phi_n\right]\delta^D. \tag{4.3}$$

To understand this expression for $\mu^B(0)$, note that there are two ways to enter default from one period to the next. First, non-excluded agents may default. Second, excluded agent can be forgiven at the start of the previous period but then default immediately. But

$$\sum_{n=0}^{+\infty} \mu^B(n)\phi_n = \mu^B(0)\left[\phi_0 + (1-\phi_0)\phi_1 + (1-\phi_0)(1-\phi_1)\phi_2 + \ldots\right] = \mu^B(0), \tag{4.4}$$

where we have used the fact that the bracketed expression adds to one if and only if

$$\sum_{n=0}^{+\infty} \Pi_{s=0}^{n}(1-\phi_n) < +\infty.$$

This should once again be intuitive. The mass of agents who exit exclusion must equal the mass of agents who enter exclusion in any stationary equilibrium. Combining expressions (4.4) and (4.3) gives the lemma. □

Since the sum of all types is one, we need

$$\mu^G + \mu^B(0)(1+\zeta) = 1$$

or, given the lemma we just established,

$$\mu^G + \mu^G(1+\zeta)\frac{\delta^D}{1-\delta^D} = 1 \Longleftrightarrow \mu^G = \frac{1}{1+\frac{\delta^D}{1-\delta^D}(1+\zeta)}.$$

In particular, maximizing the average welfare of the excluded – i.e. minimizing ζ – also maximizes the mass of active investors and hence the volume of transactions in stationary equilibrium. Unambiguously then, a policy of full but finite exclusion maximizes stationary equilibrium welfare. Theorem 4.1 – our main result – collects these results.

5 Extensions

This section discusses the robustness of our main theorem to natural variations on our benchmark environment. Appendix 9.3 studies a more complex extension in which efficiency is defined via date zero welfare rather than stationary welfare.

5.1 Lender control

Assume that lenders initiate offers in matches instead of investors. Matched investors and lenders continue to play an ultimatum game. Lenders now solve a different version of problem *(P1)*:

$$(P1') \max_{k,m} \pi(1-F(m))m - kR$$

subject to

$$\begin{aligned} k &\in K, \\ k &\leq M, \\ m &\leq Ak_t \\ \pi\left(Ak - \int_0^m \tau dF + \int_m^{\bar{\tau}} m dF\right) &\geq 0. \end{aligned}$$

so that the objective function and the participation constraint are effectively flipped. A key feature of this altered problem is that the investors' participation constraint cannot bind. Indeed,

$$\int_0^m \tau dF + \int_m^{\bar{\tau}} m dF \leq m$$

so that

$$m \leq Ak$$

implies that

$$\pi\left(Ak - \int_0^m \tau dF + \int_m^{\bar{\tau}} m dF\right) \geq 0.$$

Further, standard duality results imply that a solution to $(P1')$ exists if and only if a solution to $(P1)$ does. In particular, if parameters are such that investors are excluded in the previous version of our model, the same is true in this version. On the other hand, if a positive solution to *(P1')* exist, it must leave positive surplus on the table for investors, strictly positive surplus in fact when F has full support on $[0, \bar{\tau}]$. As for investors in good standing, equilibrium now requires that the lender's proposal (k_G, m_G) in a match with an investor in good standing maximizes

$$\pi\left(1 - F\left(m - \beta\left[V^G - V^B(0)\right]\right)\right) m - kR$$

subject to:

$$\begin{aligned} k &\in K, \\ k &\leq M, \\ m &\leq Ak, \\ \pi\left(Ak - \int_0^{m-\beta\left[V^G - V^B(0)\right]} \tau dF + \int_{m-\beta\left[V^G - V^B(0)\right]}^{\bar{\tau}} m dF\right) &\geq 0. \end{aligned}$$

Once again, the participation constraint cannot bind at any solution. Now, letting $V^L(B)$ and $V^L(G)$ denote the lender's surplus in matches with investors in bad and good standing, respectively, stationary welfare becomes:

$$\mu^G V^G + \sum_{s=0}^{+\infty} \mu^B(s) V^B(s) + \mu^G V^L(G) + \sum_{s=0}^{+\infty} \mu^B(s) V^L(B)$$

In section 2.1, we showed that front-loading punishment maximizes the average welfare of the investors in bad standing but also maximizes μ^G. Therefore, as long as $V^L(B) \leq V^L(G)$, which holds for instance in the important parametric case where investors in bad standing are excluded from loan markets, forgiveness policies that maximize $\mu^G V^G + \sum_{s=0}^{+\infty} \mu^B(s) V^B(s)$ conditional on a given level for $V^B(0)$ also maximize the objective above. As a result, our results on the efficient shape of forgiveness policies are completely unchanged.

5.2 Dead-weight loss from default

In the environment we have considered, investors bear the full direct cost τ of default. In practice default may cause resources to be wasted by parties other than the defaulting agent. Assume for instance that default carries a social cost C (in resources or in aggregate utility) that rises with the frequency of default. This makes steady state welfare equal to:

$$\mu^G V^G + \sum_{s=0}^{+\infty} \mu^B(s) V^B(s) - C \left[\mu^G V^G + \sum_{s=0}^{+\infty} \mu^B(s) V^B(s) \right] \delta^D.$$

Indeed, the mass of default in stationary equilibrium is $\left[\mu^G V^G + \sum_{s=0}^{+\infty} \mu^B(s) V^B(s)\right] \delta^D$. Since we established in section 4 that

$$\mu^B(0) = \left[\mu^G V^G + \sum_{s=0}^{+\infty} \mu^B(s) V^B(s) \right] \delta^D$$

in any stationary equilibrium, this welfare objective may be written more simply as,

$$\mu^G V^G + \sum_{s=0}^{+\infty} \mu^B(s) V^B(s) - C \mu^B(0)$$

or, given lemma 4.2,

$$\mu^G V^G + \sum_{s=0}^{+\infty} \mu^B(s) V^B(s) - \mu^G \frac{\delta^D C}{1-\delta^D} = \mu^G \left(V^G - \frac{\delta^D C}{1-\delta^D} \right) + \sum_{s=0}^{+\infty} \mu^B(s) V^B(s).$$

It follows that as long as $V^G - \frac{\delta^D C}{1-\delta^D} \geq V^B(s)$, which must hold if C or δ^D are sufficiently small at the SSE under consideration, our results hold and can be established without any change in the arguments we have made. But if C is large, the inequality may fail in which case it becomes optimal to discourage investment as much as possible given the dead-weight losses associated with default. In that case, permanent exclusion – or any forgiveness policy which asymptotically drives investment to zero – is optimal.

Finally, it may also be the case that lenders experience a private cost when default occurs. Introducing such a cost can be done simply by making it part of the lender's participation constraint in the maximization problem solved by investors when they select an offer and does not affect our main results in any way.

5.3 Private information

Suppose investors privately draw τ from the distribution $F(\cdot)$ each period before they offer a contract to lenders. Consider first the case where draws are independent across periods. For simplicity, maintain the assumption that $V^G \geq V^B(0)$ in any SSE which holds true for instance when parameters are such that investors are fully excluded from getting a loan. Because τ is known before offers are made, SSEs feature potentially separate offers $(k(\tau), (m(\tau))$ for each type τ. Separating offers from investors in good standing must solve for each particular type $\tau \in [0, \bar{\tau}]$:

$$\max \pi \left(Ak - m + \beta V^G\right) + (1-\pi)\beta V^B(0)$$

subject to

$$\begin{aligned} k &\in K \\ k &\leq M \\ m &\leq Ak \\ kR &\leq \pi m, \end{aligned}$$

subject to the no-default constraint

$$-m + \beta V^G \geq -\tau + \beta V^B(0),$$

and subject to the truth-telling constraint

$$Ak - m + \beta V^G \geq \max_{\tau' \neq \tau} \left(Ak(\tau') - m(\tau') + \beta V^G, Ak(\tau') - \tau + \beta V^B(0)\right).$$

The truth telling constraint says that investors must be better off revealing their type than mimicking the contract offered by a different type and strategically defaulting. Because τ becomes known by all when contracts are separating, each offer must be such that the payment is made with probability one when the project succeeds, which is captured by the no-default constraint. It follows that $m(\tau) = \frac{k(\tau)R}{\pi}$ for each type $\tau \in [0, \bar{\tau}]$, so that k and m are co-linear. It then follows that raising $m(\tau)$ when it is feasible improves type $\tau's$ objective.

Furthermore, increasing m weakens the truth-telling constraint since k increases linearly with m while the right-hand side of the constraint is fixed. So either the no-default constraint binds or $k(\tau) = \max K$. It follows that there is a threshold default type $\underline{\tau}$ above which investors operate at maximum scale and earn the highest surplus possible when in good standing.

At $\tau = \underline{\tau}$, the no-default constraint is exactly binding. So consider investors whose $\tau < \underline{\tau}$. They can pretend to be type $\underline{\tau}$ and default when a payment comes due. For those people,

$$Ak(\underline{\tau}) - \tau + \beta V^B > Ak(\underline{\tau}) - \underline{\tau} + \beta V^B = Ak(\underline{\tau}) - m(\underline{\tau}) + \beta V^G$$

and it follows that any contract compatible with truth-telling must give them the best feasible contract. In other words, there is no way to dissuade investors with low default cost from pretending to be investors with $\tau = \underline{\tau}$. It follows that all equilibrium offers must be such that $(k(\tau), (m(\tau))$ are independent of τ and all SSEs must therefore feature pooling. As a consequence, our results about the efficient shape of exclusion policies are fully robust to this variation.

The case with correlated draws across periods is more difficult. In that case, investors' payment histories contain information about types and offers will depend on beliefs about this type, i.e. the reputation of investors given their history, as discussed for instance by Corbae et al. (2016). Efficient forgiveness policies may likewise be belief-dependent. To see this, take an extreme example with permanent types. Type I investors have a low cost of default and may default both for strategic and non-strategic default while type II investors have a prohibitively high default cost hence only default for non-strategic reasons. Over time, type II's frequency of default must become ever closer to π so that their type becomes known with ever more precision. Since those agents never default strategically it makes no sense to punish default on part of those agents with exclusion. So, eventually, we must converge to an economy where type I is identified with arbitrary precision and may be excluded following default whereas type II investors are also identified and are not excluded following default. This stark example implies that, in general, the efficient forgiveness policy may be reputation and history dependent.

But this does not imply that the shape of exclusion policies will change. To take an even more extreme example, assume that types are known with certainty from date $t = 0$. Then type II is never excluded at all – a degenerate case of the shape we have identified as efficient. As for type I investors, can view them as comprising a sub-economy or our environment and the shape of efficient policies for them must once again have the shape we have identified as efficient in this paper. Forgiveness policies are type (reputation) dependent but, conditional on type, they must have the shape we have shown to be efficient. Generalizing these ideas to types that are persistent but not permanent will likely require new arguments, but the economics behind our main results should survive this generalization.

5.4 Long-term contracts

The random matching framework we have adopted implies that investors and lenders in a given match play a one-shot game with one another and, furthermore, that the only information each have about past actions is contained in the public record. Assume alternatively that investors and lenders have the ability to write long-term contracts and that matches may last more than one period. To simplify matters, consider the case when $K = \{I\}$ where $I > 0$ is an investment that must be made once and for all in the first period of the match. As in De Marzo and Fishman (2007), lenders can commit to contracts that specify transfers from the investor to the lender as well as a termination decision (a decision to break the match) for every possible history of output reports. Denote the termination value for lenders by $V_T^L \geq 0$, an object which captures for instance the salvage value of the installed capital. The termination payoff for investors is $V_T^I \geq 0$ and reflects for instance the value of entering into a new match. Taking termination payoffs as given, De Marzo and Fishman (2007) show that the termination option is invoked with positive probability following a sequence of bad reports. This is the case even though under the optimal contract it is understood by both parties that all reports are truthful. Threatening the borrower with termination – even when it is ex-post inefficient – is the optimal way to provide incentives against strategic default in this environment.

Now add to this environment exogenous match breaks which may occur say because the investment opportunity ended exogenously or, equivalently, because all projects have random but almost surely finite lives. The public record would then optimally consist of a bad standing signal which is activated when and only when a termination for cause occurs. In that world, an equilibrium with exogenous termination value V_G higher than termination value V_B could be sustained as a sequential equilibrium in which actions are long-term contracts rather than one-period decisions. While establishing this is beyond the scope of this current paper, it seems likely to us that equilibrium in which investors are temporarily excluded following termination for cause can be built and that it remains efficient for termination to be front-loaded. Because the long-term contractual environment is much more challenging than the random matching framework we have employed, we recognize that formalizing these ideas will be complex.

5.5 Risk-aversion

Much of the literature on endogenous default – Corbae et al. (2005) and Kehoe and Levine (1993), for two prominent examples – focuses on the relationship between exclusion threats and the endogenous level of risk-sharing. This subsection considers a version of our environment with a risk-sharing motive and shows that the optimal shape of the exclusion policy is unchanged.

Assume that investors have time separable preferences with the same discount rate as before but a Von Neumann-Morgenstern period utility function U that is concave with $|U(0)| < +\infty$. Further assume that the default cost τ is measured in consumption equivalent units. Since lenders are risk-neutral, investors and lenders now have an incentive to share risk. In principle, loans could therefore stipulate a positive transfer from lenders to the investor when the project fails. Under the premise that lender actions are not recorded however, lenders have no incentive to make such a transfer ex-post and so we will maintain the restriction that transfers are zero when the project fails. As has been the case throughout this paper, we continue to focus on stationary and symmetric strategies. Any SSE (k_B, m_B, k_G, m_G) is associated with continuation values V^G and $V^B(n)$ for investors in each possible state that solve

$$V^G = (1-\pi)\beta V^B(0) + \pi E_\tau \max\left\{U\left(Ak_G - m_G\right) + \beta V^G, U\left(Ak_G - \tau\right) + \beta V^B(0)\right\} \quad (5.1)$$

and, for investors who have been in bad standing for $n \geq 0$ periods,

$$V^B(n) = \phi_n V^G + (1-\phi_n)\left[\pi E \max\left\{U\left(Ak_B - m_B\right), U\left(Ak_B - \tau\right)\right\} + \beta V^B(n+1)\right]. \quad (5.2)$$

These continuation values remain uniquely defined at any SSE. Investors in good standing choose to pay if the project succeeds and once they discover their default cost τ if

$$U\left(Ak_G - m_G\right) + \beta V^G \geq U\left(Ak_G - \tau\right) + \beta V^B(0) \quad (5.3)$$

which defines a unique threshold $\tau\left(k, m_G, V^G, V^B(0)\right)$ such that investors pay if and only if their ex-post default cost realization is above that threshold. Lenders break even provided:

$$k_G R \leq \pi \left(1 - F\left(\tau\left(k, m_G, V^G, V^B(0)\right)\right)\right) m_G. \tag{5.4}$$

Likewise, problem $(P1)$ needs to be amended to:

$$(P1'') \max_{k,m} \pi U \left(Ak - \int_0^m \tau dF + \int_m^{\bar{\tau}} m dF\right)$$

subject to:

$$\begin{aligned} k &\in K, \\ k &\leq M, \\ m &\leq Ak \\ kR &\leq \pi(1 - F(m))m. \end{aligned}$$

With this notation in place, a symmetric and stationary equilibrium (SSE) can once again be defined. An SSE is a quadruple (k_B, m_B, k_G, m_G), the associated policy functions V^G and $V^B(n)$, and stationary acceptance policies for lenders for every possible match and every possible offer, such that

1. The acceptance policy is optimal for lenders for every possible offer in any match. That is, offers are accepted if and only if they satisfy 5.4 when made by investors in good standing and 3.4 when made by investors in bad standing;
2. Proposal (k_B, m_B) solves problem $(P1'')$;
3. Proposal (k_G, m_G) maximize

$$\pi U \left(Ak - \int_0^{m-\beta\left[V^G - V^B(0)\right]_t} \tau dF + \int_{m-\beta[V^G - V^B(0)]}^{\bar{\tau}} m dF\right)$$

subject to:

$$\begin{aligned}
k &\in K, \\
k &\leq M, \\
m &\leq Ak, \\
\pi(1 - F\left(\tau\left(k, m, V^G, V^B(0)\right)\right))m &\geq kR.
\end{aligned}$$

Our existence arguments go through virtually unchanged and lemma 3.4 continues to hold so that the same definition of efficiency can be stated and justified. With the shortcut notation

$$U_B \equiv \pi E \max\left\{U\left(Ak_B - m_B\right), U\left(Ak_B - \tau\right)\right\}$$

we can write

$$\begin{aligned}
V^B(0) &= \phi_0 V^G + (1-\phi_0)\left[U_B + \phi_1 \beta V^G\right] + (1-\phi_0)(1-\phi_1)\left[\beta U_B + \phi_2 \beta^2 V^G\right] + \ldots \\
&= \frac{U_B}{1-\beta} + \phi_0 \widetilde{V}^G + (1-\phi_0)\phi_1 \beta \widetilde{V}^G + (1-\phi_0)(1-\phi_1)\phi_2 \beta^2 \widetilde{V}^G + \ldots
\end{aligned}$$

where

$$\widetilde{V}^G = V^G - \frac{U_B}{1-\beta}.$$

A similar expression obtains for $V^B(n)$ for all $n > 0$. In other words, introducing general preferences results in level shift of continuation utility functions – which cannot have any effect on welfare rankings of exclusion policy function - and a parallel shift of the gains from forgiveness given V^G. It follows that the same exclusion policy as in the linear case continues to maximize the average welfare of investors in bad standing. The fact that this also maximizes the volume of transactions follows from the same arguments as before.

5.6 Self-financing

Investors enter each period without any assets or endowments which simplifies the analysis by implying that the entire capital investment must be financed by lenders. Assume instead that investors are endowed with a quantity $a \geq 0$ of the good at the start of each period. In the version of our model with time separable, Von Neumann-Morgenstern preferences, problem

$(P1)$ becomes:

$$\max_{k,m,s} \pi U \left(Ak - \int_0^m \tau dF + \int_m^{\bar{\tau}} mdF + (a-s)R \right) + (1-\pi)U\left((a-s)R\right)$$

subject to:

$$\begin{aligned} k+s &\in K, \\ k &\leq M, \\ m &\leq A(k+s) \\ kR &\leq \pi(1-F(m))m \end{aligned}$$

where $s \geq 0$ is the investor's contribution to the project which can be implemented either via a deposit over which the lender has a lien or a direct investment. The problem for investors in good standing can be similarly adjusted. Saving and self-financing some of the capital they employ enables investors to increase the size of the feasible offer set. SSEs continue to be fully summarized by a quadruple (k_B, m_B, k_G, m_G) and, although the self-investment choice may vary across different equilibria, lemma 3.4 continues to hold, and our efficiency arguments can be applied virtually unchanged to lead to the same optimal shape of forgiveness policy. Put another way, allowing for self-financing via fixed investor endowments has no impact on our main results. With exogenous but stochastic endowments, different investors may start each period with different endowments and hence may make different financing choices. If we maintain the assumption that forgiveness policies are independent of the size of the loan on which investors defaulted, continuation values remain common across investors and our arguments will once again go through essentially unaffected.

The case in which investor assets are the result of endogenous saving decisions is substantially more difficult.[12] When goods are storable and investors make saving decisions from period to period – absent an equalization trick à la Lagos and Wright (2005) – any stationary equilibrium involves a non-degenerate distribution of assets where an investor's wealth is a function of the agent's entire income history. The shape and length of the forgiveness policy must now impact saving policies. In fact, forgiveness policies could in principle depend on the assets of agents in default or, equivalently, on the size of the loan on which agents defaulted.

[12] For that reason, much of the literature inspired by Kehoe and Levine (1993) and Alvarez and Jerman (2000) assumes as we do that goods are non-storable.

While we do not have obvious reasons to think that this affects the fact that at least some front-loading of exclusion is optimal, we leave a complete treatment of this significantly more complex case for future work.[13]

5.7 Partial exclusion

In section 4 we restricted parameters so that $(k_B, m_B) = (0, 0)$ in all SSEs and investors in bad standing do not get a loan. While doing so economizes on notation and makes for a tight mapping between our model and traditional models of endogenous default in which exclusion is usually assumed to be full, our environment makes it easy to study situations where $(k_B, m_B) > (0, 0)$. The only condition we need for our results to go through is that $V^B(n) \leq V^G$ for all n in SSE, so that being in bad standing is in fact detrimental.[14]

In practice, borrowers with recent defaults on their record – borrowers who recently filed for bankruptcy, e.g., do have access – albeit comparatively limited – to some credit, as documented for instance by Han and Li (2011). In fact the evidence suggests that access to credit improves over time while households are under a bankruptcy flag, in part because judges may become more lenient in authorizing new borrowing when distance from the most recent bankruptcy has increased. Our model can accommodate this feature since it also allows for equilibrium where lenders condition their actions on the time since borrowers have defaulted. As long as the evolution of an investor standing – forgiveness odds in particular – are exogenous, (k_B, m_B) must solve $(P1')$. Because solutions to $(P1')$ may be multiple, one could construct equilibrium with different loans for investors in bad standing, but this multiplicity does not provide fundamental reasons why access to credit should improve over time. A version of our model where investors gradually improve their standing as they make payments while under default, on the other hand, does provide economic forces that could sustain improving terms over time for those investors.

[13] Modeling endogenous savings would also enable one to revisit the question posed by Bulow and Rogoff (1989) and revisited by Hellwig and Lorenzoni (2009) of whether excluding borrowers from saving following default is necessary to sustain credit in environments such as ours where commitment is limited.

[14] Because in environments like ours the set of equilibrium is large and can contain surprising outcomes, we cannot rule out the opposite situation in full generality. At the same time, as we have discussed in various examples already, it is easy to restrict parameters so that equilibrium with $V^B(n) \leq V^G$ for all n do exist.

6 Exclusion length

In our model, efficient forgiveness policies must be deterministic and punishment must be finite in length. But what should that length be? In this section, we show that optimal exclusion length depends in non-monotonic ways on fundamental parameters. For concreteness, we continue to concentrate on SSEs in which exclusion is complete for investors in bad standing. In that case and given a forgiveness policy,

$$V^B(0) = \phi_0 V^G + (1-\phi_0)\phi_1\beta V^G + (1-\phi_0)(1-\phi_1)\phi_2\beta^2 V^G + \ldots = \kappa V^G$$

in all SSEs where $\kappa = \frac{V^B(0)}{V^G} < 1$ is an exclusion discount of sorts. Furthermore, given that all efficient forgiveness policies are deterministic and finite κ uniquely pins down and is negatively related to the length of exclusion associated with any efficient forgiveness policy. What is more, the participation constraint of lenders may be rewritten as:

$$\pi\left(1 - F\left(m_G - \beta(1-\kappa)V^G\right)\right) m_G \geq k_G R. \tag{6.1}$$

Note that under the parametric premise that there is no solution to $(P1)$ so that $(k_B, m_B) = (0,0)$ in all SSEs, (k_G, m_G) implies V^G given κ. The search for the optimal level of κ may then be organized as follows. Given κ,find a solution (k_G, m_G) to condition 6.1 that makes V^G highest. Since it makes V^G highest it also maximizes $V^B(n)$ for all n. We can then update κ until average steady state welfare is highest.

To illustrate these ideas, we will begin with a simple example. Consider a version of our economy where the distribution F of default costs only puts mass at one point τ and where $K = \{1\}$ so that all projects must be operated at a unit scale when they are activated. In particular, $k_G = 1$ in all SSEs. If $\tau \geq \frac{R}{\pi}$, lenders can set $m = \frac{R}{\pi}$, which exposes them to zero strategic default and enables them to break even. In that case, the optimal length of exclusion is zero.

When $\tau < \frac{R}{\pi}$, on the other hand, investors do have static incentives to default for strategic reasons. Absent some exclusion threat, they would all default which cannot be in equilibrium. This yields a key simplification in the construction of equilibrium. Because investors are homogenous ex-ante, the optimal solution has to be such that all investors pay as long as they can. As a result, in this simple example, the optimal exclusion policy must be the one that makes payment among investors with successful projects just optimal.

Investors all pay their loans when

$$m \leq (1-\kappa)\beta V^G + \tau. \tag{6.2}$$

It is clearly optimal to make this inequality tight. Otherwise investors are excluded for longer than is strictly needed. Furthermore, it is not just incentive compatible but also optimal for investors to pay since this makes $\mu^G = 1$. The break-even condition for lenders becomes:

$$\pi m = \pi((1-\kappa)\beta V^G + \tau) = R, \tag{6.3}$$

and we can solve for V^G as follows:

$$V^G = \pi\left(A - m + \beta V^G\right) + (1-\pi)\beta\kappa V^G$$

since, in this case, investors only default for non-strategic reasons. This equation together with the fact that

$$m = (1-\kappa)\beta V^G + \tau = \frac{R}{\pi}$$

enables us to solve out for V^G and obtain a condition which the exclusion discount κ must solve in any equilibrium:

$$\frac{R - \pi\tau}{\pi\beta(1-\kappa)} = \pi A - R + \frac{R - \pi\tau}{(1-\kappa)} + \frac{(1-\pi)\kappa(R-\pi\tau)}{\pi(1-\kappa)}.$$

A bit of algebra then yields a closed-form expression for the optimal exclusion policy:

$$\kappa = \frac{\pi^2(A-\tau) - \frac{1}{\beta}(R - \tau\pi)}{\pi^2(A-\tau) - (R-\tau\pi)} \tag{6.4}$$

It follows that when default costs are homogenous at given value τ, the optimal exclusion discount solves

$$\kappa = max\left(\frac{\pi^2(A-\tau) - \frac{1}{\beta}(R - \tau\pi)}{\pi^2(A-\tau) - (R-\tau\pi)}, 0\right).$$

In this simple example, exclusion length falls with investor patience (β), the project payoff (A), project quality (π), and with the direct punishment (τ) associated with default. An increase of patience increases the dissuasive power of exclusion and makes it possible to shorten its length. An increase in project payoff (y) makes exclusion more costly since the benefit from investing are higher. Raising π, likewise, makes the value of participation higher. While, at the same time, it does lead to more opportunities to default for strategic reasons, the positive effect on the value of participation dominates. Probably most interesting in this example is the relationship between incentives to default for strategic defaults and exclusion length. When the project succeeds, the cost of defaulting is two-fold: exclusion and direct punishment τ. When τ is higher, exclusion becomes less useful and the planner can shorten it which leads to a higher volume of transactions. On a basic level, the fact that the propensity to default for strategic reasons matters for optimal exclusion length is not surprising. After all, the only point of exclusion is to discourage strategic default.

While this simple example produces sharp comparative statics, a simple variation shows that different versions of our economy can yield very different predictions. Indeed, assume that $K = \{k_L, k_H\}$ where $k_L < k_H$ so that the project may now be operated either on small scale or on a large scale. Assuming that problem $(P1)$ remains such that $(k_B, m_B) = (0, 0)$, there are two possible types of SSEs with positive investment depending on whether k_G equals k_L or k_H. In either case, we must still have

$$(1-\kappa)\beta V^G + \tau = \frac{k_G R}{\pi}.$$

Furthermore, the structure of our environment is such that V^G is linear in k_G, holding other parameters the same. The equation above implies that if both types of SSEs exist, κ must be higher when $k_G = k_H$ than when $k_G = k_L$. This is intuitive. When more funds have to be committed to the project by lenders, all else equal, more incentives to repay must be provided to investors.

The crux of the issue for our purposes is that k_G is endogenous and the fundamental factors that cause capital choices to go up must have direct effects on incentives to repay as well. To understand this, assume that π is initially such that in the SSE where$k_G = k_L$ is feasible. One way to restrict parameters so that this holds is to set make β low enough so that

$$\beta V^G + \tau < \frac{k_H R}{\pi}$$

so that, even if $\kappa = 0$ (i.e. even if exclusion is permanent) incentives to repay cannot be provided if the project is operated on a large scale but that

$$\beta V^G + \tau > \frac{k_L R}{\pi}$$

so that sufficient incentives can be provided to support $k_G = k_L$ as an SSE.

Now assume that π rises in such a way that the first inequality is reversed. This can be made to happen because $\beta V^G + \tau$ rises with π while $\frac{k_H R}{\pi}$ falls. With A or π sufficiently large or β sufficiently low, the SSE with a large capital scale must be associated with higher average welfare than the SSE with low capital. But if the individual rationality just holds at the new parameter constellation, κ must be close to zero, and the length of exclusion must increase. So we have built a case where, in exact contrast to the previous example, an increase in π causes an increase in exclusion length. The same ideas show that an increase in τ can cause large scale production to become feasible which may once again cause an increase in exclusion length: a higher default cost can be associated with longer exclusion.

Section 6 considers yet another variation of this example in which default costs are no longer homogenous. Introducing heterogeneity in default costs, once again, has ambiguous effects on the optimal length of exclusion.

The bottom line is that optimal length of exclusion is a complex function of all fundamental parameters in our model. Whereas the optimal shape of exclusion policies can be uniquely pinned down as we have shown in this paper, the impact of parameters on the optimal length of exclusion can only be established on a case-by-case basis. Outside of the simple class of examples we have considered in this section, quantitative explorations are bound to be necessary to pin down the relationship between the optimal length of exclusion and fundamental parameters.

7 Conclusion

In a canonical model of borrowing and lending with endogenous default, providing public records that make exclusion following default sustainable in equilibrium enlarges the set of equilibrium that can be supported. We show that, efficiently from the point of view of long-term welfare, the public record technology should be designed so as to imply exclusion for a finite and deterministic number of periods following default. Not only does front-loading exclusion maximize the volume of transactions, it also maximizes the welfare of the excluded, conditional on the fact that some punishment generally needs to be imposed to sustain positive lending. This set of results is robust to a host of considerations but they remain to be generalized to environments in which investors accumulate resources over time and to environments with long-term contracts. Our view is that the simple economics behind our results should generalize to those settings although making that case will require substantially different arguments from ours.

8 References

Azariadis, C., and L. Kass (2013), "Endogenous credit limits with small default costs," *Journal of Economic Theory,* 148, pp. 806-824.

Alvarez, F., and Jermann, U. J. (2000), "Efficiency, Equilibrium, and Asset Pricing with Risk of Default," *Econometrica*, 68, pp. 775–797.

Araujo, A., Fajardo, J., and M. R. Pascoa (2005), "Endogenous collateral," *Journal of Mathematical Economics*, 41, pp. 439-462.

Bethune, Z., Rocheteau, G., and T-W Hu (2018), "Indeterminacy in Credit Economies," *Journal of Economic Theory*, 175, 556–584.

Bolton, P., and D. S. Scharfstein (1990), "A Theory of Predation Based on Agency Problems in Financial Contracting," *American Economic Review*, 80 (1), 93-106.

Bond, P. and A. Krishnamurthy (2004), "Regulating Exclusion from Financial Markets," *Review of Economic Studies*, 71, pp. 681-707.

Bulow, J., and K. Rogoff (1989), "Sovereign Debt: Is to Forgive to Forget?", *American Economic Review*, 79, 43–50.

Corbae, D., Chatterjee, S., Dempsey, K., and J. V. Rios-Rull (2016), "A Theory of Credit Scoring and the Competitive Pricing of Default Risk," Working Paper, University of Wisconsin – Madison.

DeMarzo, P. and M. Fishman (2007) "Optimal Long-Term Financial Contracting," *Review of Financial Studies*, 20, pp. 2079–2128.

Dubey, P., Geanakoplos, J., and M. Shubik (2005), "Default and Punishment in General Equilibrium," *Econometrica*, 73, pp. 1468-0262.

Chatterjee, S., Corbae, D., Nakajima, M., and V. Rios-Rull (2007), "A Quantitative Theory of Unsecured Consumer Credit with Risk of Default," *Econometrica*, 75, pp.1525-1589.

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

Foote, C., Gerardi, K., and P. Willen (2008), "Negative Equity and Foreclosure: Theory and Evidence," *Journal of Urban Economics*, 64, pp. 234-45.

Gerardi, K., Herkenhoff, K., Ohanian, L., and P. Willen (2015), "Can't Pay or Won't Pay? Unemployment, Negative Equity, and Strategic Default," Research Department Working Paper 15-13, Federal Reserve Bank of Boston.

Ghent, A. C., and M. Kudlyak (2011), "Recourse and Residential Mortgage Default: Theory and Evidence from U.S. States," *Review of Financial Studies,* 24, pp. 1814-47.

Green, E. and R. Porter (1984), “Noncooperative Collusion under Imperfect Price Information,” *Econometrica*, 52(1), pp. 87-100.

Gu, C., F. Mattesini, C. Monnet, and R. Wright (2013a), “Banking: A New Monetarist approach,” *Review of Economic Studies,* 80, pp. 636-662.

Gu, C., F. Mattesini, C. Monnet, and R. Wright (2013b), “Endogenous credit cycles,” *Journal of Political Economy,* 121, pp. 940 - 965.

Guiso, L., Zingales, L., and P. Sapienza (2013), “The Determinants of Attitudes towards Strategic Default on Mortgages,” *Journal of Finance,* 68, pp. 1473–1515.

Han, S., and G. Li (2011), “Household Borrowing after Personal Bankruptcy,” *Journal of Money, Credit and Banking,* 43, pp 491–517.

Hellwig, C., and G. Lorenzoni (2009), “Bubbles and Self-Enforcing Debt,” *Econometrica*, 77, 1137–1164.

Kehoe, T., and D. Levine (1993), “Debt Constrained Asset Markets,” *Review of Economic Studies*, 60, pp. 865–88.

Kehoe, T., and D. Levine (2001), “Liquidity Constrained Markets versus Debt Constrained Markets,” *Econometrica*, 69, 575–98.

Kehoe, T., and D. Levine (2008), “Bankruptcy and Collateral in Debt Constrained Markets,” in Roger E. A. Farmer, editor, *Macroeconomics in the Small and the Large: Essays on Microfoundations, Macroeconomic Applications*, and Economic History in Honor of Axel Leijonhufvud, Edward Elgar.

Kocherlakota, N., and N. Wallace (1998), “Incomplete Record-Keeping and Optimal Payment Arrangements,” *Journal of Economic Theory*, 81, 272–289.

Kirpalani, R. (2018), “Endogenously Incomplete Markets with Equilibrium Default,” Working Paper, Pennsylvania State University.

Krueger, D. and F. Perri (2005), “Does Income Inequality Lead to Consumption Inequality? Evidence and Theory,” *Review of Economic Studies*, 73, pp. 163-193

Livshits, I., MacGee, J., and M. Tertilt (2007), “Consumer Bankruptcy - A Fresh Start,” *American Economic Review*, 97, 402-18.

Lagos, R. and R. Wright (2005), “A unified framework for monetary theory and policy,” *Journal of Political Economy*, 113(3), 463-484

Liu, Q. and A. Skrzypacz (2013), “Limited records and reputation bubbles,” *Journal of Economic Theory*, 2014, 151, pp. 2–29.

Quintin, E. (2013), “On Existence in Equilibrium Models with Endogenous Default,” *Jour-*

nal of Mathematical Economics, 49, pp. 418–421.

Quintin, E (2012), "More Punishment, Less Default?," *Annals of Finance*, 8, 2012, pp. 427-454.

Tierney, L. (1994), "Markov chains for exploring posterior distributions," *Annals of Statistics,* 22, pp 1701–1762.

9 Appendix

9.1 Simplification of expression $\mathcal{P}$

We need to show that

$$(1-\phi_0)V^B(1)+(1-\phi_0)(1-\phi_1)V^B(2)+...$$

is constant given constraint 4.2. To see it, write the expression as follows:

$$\begin{array}{rll}
(1-\phi_0)\phi_1 V^G & +(1-\phi_0)(1-\phi_1)\phi_2\beta V^G & +(1-\phi_0)(1-\phi_1)(1-\phi_2)\phi_3\beta^2 V^G+... \\
(1-\phi_0)(1-\phi_1)\phi_2 V^G & +(1-\phi_0)(1-\phi_1)(1-\phi_2)\phi_3\beta V^G & +(1-\phi_0)(1-\phi_1)(1-\phi_2)(1-\phi_3)\phi_4\beta^2 V^G+... \\
(1-\phi_0)(1-\phi_1)(1-\phi_2)\phi_3 V^G & +(1-\phi_0)(1-\phi_1)(1-\phi_2)(1-\phi_3)\phi_4\beta V^G & +(1-\phi_0)(1-\phi_1)(1-\phi_2)(1-\phi_3)(1-\phi_4)\phi_5\beta^2 V^G \\
+.... & &
\end{array}$$

Now sum the whole infinite expression column by column. The coefficients in the first column sum up to

$$(1-\phi_0)-\Pi_{n=0}^{+\infty}(1-\phi_n).$$

Given that $\sum_{n=0}^{+\infty}\Pi_{s=0}^{n}(1-\phi_n)<+\infty$, the second term is zero.[15] As a result, the first column gives $(1-\phi_0)V^G$. For the second column sum all weights and apply the same argument as above to get $(1-\phi_0)(1-\phi_1)$. For the third column the sum of all weights is $(1-\phi_0)(1-\phi_1)(1-\phi_2)$. So summing it all we get:

$$(1-\phi_0)V^G+(1-\phi_0)(1-\phi_1)\beta V^G+(1-\phi_0)(1-\phi_1)(1-\phi_2)\beta^2 V^G+....$$

Now note that each term has a $1-\phi_0$ factor so that it can each be split into two pieces to get:

$$V^G+(1-\phi_0)\beta V^G+(1-\phi_0)(1-\phi_1)\beta^2 V^G+...-\phi_0 V^G-(1-\phi_0)\phi_1\beta V^G-(1-\phi_0)(1-\phi_1)\phi_2\beta^2 V^G-...$$

The second part of the expression, given the constraint, is simply $V^B(0)$. The first part, other than for the very first term, yet again features a common factor $(1-\phi_0)$ which can be used

[15] In fact, not only is $\Pi_{n=0}^{+\infty}(1-\phi_n)=0$ but, with the convention that if $\phi_n=1$ then $\phi_s=1$ for all $s=1$, then $\Pi_{n=s}^{+\infty}(1-\phi_n)=0$ for all $s\geq 0$. This means that the reasoning we apply to the first column applies similarly to all other columns. The convention can be imposed without any loss of generality since $\phi_n=1$ caps exclusion at n periods with probability one.

to split it into two subparts, leaving us with:

$$V^G - V^B(0) + \beta V^G + (1-\phi_0)\beta^2 V^G + \ldots$$

$$-\beta\left\{\phi_0 V^G - (1-\phi_0)\phi_1\beta V^G - (1-\phi_0)(1-\phi_1)\phi_2\beta^2 V^G - \ldots\right\}$$

By the constraint, the final line in this expression is nothing but $\beta V^B(0)$. Continuing in this fashion shows that the whole sum is

$$V^G - V^B(0) + \beta\left[V^G - V^B(0)\right] + \beta^2\left[V^G - V^B(0)\right] = \frac{V^G - V^B(0)}{1-\beta}$$

which is a constant given V^G and $V^B(0)$, as claimed.

9.2 Full but finite exclusion maximizes $\mathcal{P}$

Given the previous result, maximizing $\mathcal{P}$ amounts to minimizing

$$(1-\phi_0) + (1-\phi_0)(1-\phi_1) + (1-\phi_0)(1-\phi_1)(1-\phi_2) + \ldots$$

subject to the restriction that condition 4.2 must hold i.e. that $V^B(0)$ is what it needs to be to support the stationary equilibrium. Traditional variational arguments show that this is done by adopting the policy described in proposition 4.1.

To see this, assume first that $\phi_2 = 1$ so that excluded investors are sure to return to markets after two periods of exclusion. In that case, the problem boils down to

$$\min(1-\phi_0) + (1-\phi_0)(1-\phi_1)$$

subject to:

$$\phi_0 + (1-\phi_0)\phi_1\beta + (1-\phi_0)(1-\phi_1)\beta^2 = \frac{V^B(0)}{V^G}.$$

Now add and subtract $(1-\phi_0)(1-\phi_1)\beta$ to the left-hand side of the constraint to get:

$$\phi_0 + (1-\phi_0)\beta - (1-\phi_0)(1-\phi_1)(\beta-\beta^2) = \frac{V^B(0)}{V^G}.$$

The proposition holds if $\phi_0 > 0 \Longrightarrow \phi_1 = 1$. Assume by way of contradiction that $\phi_0 > 0$ but $\phi_1 < 1$. Then it is possible to reduce ϕ_0 by some $\epsilon > 0$. This causes the first two terms of the constraint to fall by a total of $\epsilon(1-\beta)$. Maintaining the constraint level thus requires that $(1-\phi_0)(1-\phi_1)$ falls by

$$\frac{\epsilon(1-\beta)}{\beta-\beta^2}=\frac{\epsilon}{\beta}.$$

But then $(1-\phi_0)$ rises by ϵ while $(1-\phi_0)(1-\phi_1)$ falls by $\frac{\epsilon}{\beta}$, which improves (i.e. lowers) the objective strictly, the contradiction we sought.

Assume now that $\phi_3 = 1$. The objective becomes

$$\min(1-\phi_0)+(1-\phi_0)(1-\phi_1)+(1-\phi_0)(1-\phi_1)(1-\phi_2)$$

subject to:

$$\phi_0+(1-\phi_0)\phi_1\beta+(1-\phi_0)(1-\phi_1)\phi_2\beta^2+(1-\phi_0)(1-\phi_1)(1-\phi_2)\beta^3=\frac{V^B(0)}{V^G}.$$

Rewrite the constraint as

$$\phi_1+(1-\phi_1)\phi_2\beta+(1-\phi_1)(1-\phi_2)\beta^2=\frac{\frac{V^B(0)}{V^G}-\phi_0}{\beta(1-\phi_0)}.$$

This makes it clear that holding ϕ_o constant the problem in ϕ_1 and ϕ_2 is exactly the same as before. This implies as before that if $\phi_1 > 0$ then $\phi_2 = 1$. But then in that case we are back once again to the problem above which implies that if $\phi_0 > 0$ then $\phi_1 = 1$. If on the other hand $\phi_1 = 0$ then the constraint reads as

$$\phi_0+(1-\phi_0)\phi_2\beta^2+(1-\phi_0)(1-\phi_2)\beta^3=\frac{V^B(0)}{V^G}.$$

But we can then invoke the same argument as above (add and subtract $(1-\phi_0)(1-\phi_2)\beta^2$ to the left-hand side of the constraint and proceed) to conclude that if $\phi_0{>}0$ and $\phi_1 = 0$ then $\phi_2 = 1$ is optimal. But we already know that if $\phi_2 = 1$ then $\phi_0{>}0$ and $\phi_1 = 0$ cannot be optimal.

All told then and proceeding recursively, the solution has to be such that if $\phi_s > 0$ for some s then $\phi_{s+1} = 1$, as long as ϕ is eventually 1. Under that premise, minimizing

$$(1-\phi_0)+(1-\phi_0)(1-\phi_1)+(1-\phi_0)(1-\phi_1)(1-\phi_2)+\dots$$

is done by selecting the unique policy $\{\phi_t^*\}_{t=0}^T$ that satisfies the conditions of proposition 4.1 and meets the punishment constraint exactly.

To complete the proof then, we only need to argue that the premise that ϕ is eventually 1 is without loss of generality. Denote by $\{\bar{\phi}_t\}_{t=0}^T$ a policy that minimizes the above objective without imposing that restriction. That policy must be such that $\bar{\phi}_t > 0$ for at least one t. So there must be a first non-zero term. Without loss of generality, assume $\bar{\phi}_0 > 0$. (The argument below can be shifted forward if $\bar{\phi}_0$ is the first non-zero term is further along the ϕ sequence.)

Fix $\epsilon > 0$. Pick T high enough so that

$$\sum_{s=T}^{+\infty} \beta^s V^G < \frac{\epsilon}{k}$$

where k is a positive constant to be specified below. This cutoff has the property that the expected value

$$\left(\bar{\phi}_0 + (1-\bar{\phi}_0)\phi_1\beta + (1-\bar{\phi}_0)(1-\bar{\phi}_1)\bar{\phi}_2\beta^2 + \ldots\right) V^G$$

accounted for by $\{\bar{\phi}_t\}_{t=0}^{T-1}$ has to be within $\frac{\epsilon}{k}$ of $V^B(0)$. Now consider the alternative policy $\left\{\hat{\phi}_t\right\}_{t=0}^T$ which coincides with $\{\bar{\phi}_t\}_{t=0}^T$ up to $T-1$ but is identically one thereafter. That policy lowers the objective vis-a-vis $\{\bar{\phi}_t\}_{t=0}^T$ but may exceed $V^B(0)$ by at most ϵ. This can be rectified by lowering $\hat{\phi}_0 = \bar{\phi}_0 > 0$ by an amount less than ϵ as long as k is selected to be large enough. The resulting policy gives an objective value within ϵ of $\{\bar{\phi}_t\}_{t=0}^T$ and this is true, a fortiori, of $\{\phi_t^*\}_{t=0}^T$. Since $\{\phi_t^*\}_{t=0}^T$ is ϵ-optimal for all ϵ, it achieves the same minimum as $\{\bar{\phi}_t\}_{t=0}^T$, and is therefore optimal. This completes the proof.

9.3 Maximizing date zero welfare

We have made the case that exclusion policies that maximize stationary welfare must front load-punishment. Does front-loading punishment also maximize the welfare of the investors who happen to be alive at date zero, with types arbitrarily distributed? If possible, it is clearly efficient to forgive all the investors in bad standing at date 0 since punishing them has no remaining impact on incentives while forgiving them makes the volume of transactions as high as it can be at date 0. We will consider both the case where those amnesties are possible and the case where they cannot.

For simplicity, we will concentrate our attention once again on the situation where problem $(P1)$ has no solution so that investors in bad standing do not operate their project. Furthermore, we also continue to restrict our attention on symmetric SSEs, defined as before, with the implied constant continuation utilities. Such an SSE continues to exist for all possible forgiveness policies and the arguments for why this is so are unchanged.

In keeping with our earlier notation, denote the mass of investors in good standing at date 0 by μ_0^G. When amnesty is an option, it is efficient to set $\mu_0^G = 1$ via immediate forgiveness but, in general, the mass $\mu_0^B(n)$ of incumbent agents who, at date 0, have been excluded for n period may be positive. Efficient forgiveness policies must maximize

$$\mu_0^G V^G + \sum_{s=0}^{+\infty} \mu_0^B(s) V^B(s)$$

where V^N is the expected lifetime utility of agents who are not excluded at date 0 under the assumed SSE while $V^B(s)$ is the same for agents who have been excluded for s periods subject to

$$V^B(0) = \phi_0 V^G + (1-\phi_0)\phi_1 \beta V^G + (1-\phi_0)(1-\phi_1)\phi_2 \beta^2 V^G + \dots$$

Importantly, we no longer need to impose $\sum_{n=0}^{+\infty} \Pi_{s=0}^{n}(1-\phi_n) < +\infty$. As we will see below in fact, policies that imply that eventually all agents are excluded may well be optimal when maximizing date zero welfare.

To proceed, assume an amnesty of all agents who are in bad standing at date 0 is feasible. Following amnesty, the economy begins date 0 will all agents in good standing. In that case, SSEs are trivial to rank in terms of welfare. The lifetime utility of non-excluded agents at date 0 is fully summarized by the value $V^B(0)$ of becoming excluded since it implies the terms on loans and, in turn, V^G which is all we need to know when μ_0^G=1. It follows that exclusion policies that imply the same $V^B(0)$ all result in the same welfare level. This implies that multiple policy shape may now be optimal. Indeed, forgiveness policies with very different profiles may imply the same level $V^B(0)$ of initial punishment. To show that this is in fact a possibility, consider a parametric version of our environment in which agents discount the future at a rate $\beta = 0.85$, the probability of success is $\pi = 0.95$, output when positive is $y = 1.5$, and the opportunity cost of funds is $R = 1.2$. The default cost τ follows a log-normal distribution with location parameter 2 and dispersion parameter 1.

Figure 1: Welfare-equivalent exclusion policies when amnesty is an option

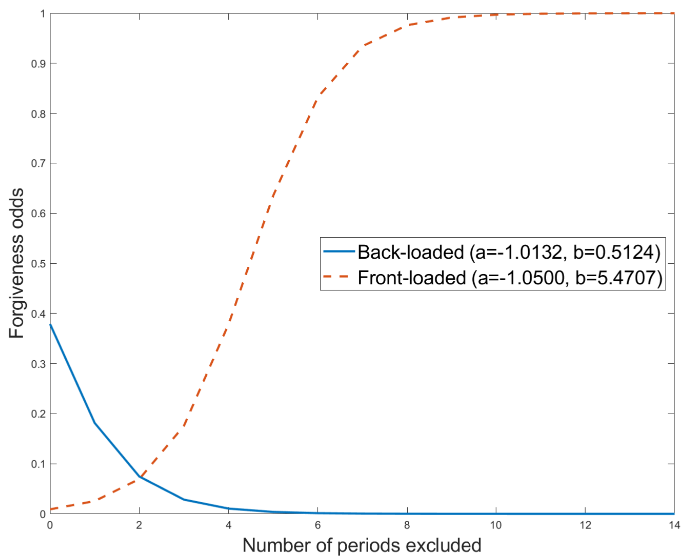

Consider then a flexible, sigmoid class of forgiveness policies characterized by two parameters (a, b) such that for $n \geq 0$:

$$\phi_n = \frac{1}{1 + \exp\left[a(n-b)\right]}$$

This family can closely approximate most monotonic exclusion policies including step-functions in which case b pins down the location of the inflection point while a pins down the steepness of the inflection. The sign of a determines whether forgiveness odds rise or fall with time in exclusion. This specification thus allows both for policies that front-load and policies that back-load punishment.

In this parametric case, the two policies displayed in figure 1 imply the same equilibrium level of $V^B(0)$ hence the same level of welfare. We located these two welfare equivalent policies by searching numerically for welfare maximizing policies given $V^B(0)$ starting from different initial conditions. The search stops the moment the procedure has found a candidate that achieves the target level of $V^B(0)$[16] Many other policies, including non-monotonic ones, deliver

[16]A numerical procedure is needed to look for those welfare-equivalent policies since any change in the policy implies a change in lending terms and value function. The first step of our procedure solves for equilibrium

the same welfare level but those two specific examples suffice to convey two key ideas. First, when amnesty is an option, the efficient shape of exclusion policies is indeterminate. One of the two policies displayed in the figure front-loads punishment in the sense that forgiveness odds are initially low but then rise, while the other back-loads punishment. In other words, while front-loading is uniquely efficient under the long-term perspective we adopted in previous sections, it is only weakly optimal when maximizing the welfare of investors alive at date 0 *and* when amnesty is an option. Second, policies that violate ergodicity criterion $\sum_{n=0}^{+\infty} \Pi_{s=0}^{n}(1-\phi_n) < +\infty$ can be efficient. The back-loaded policy shown in the figure has forgiveness policies converge to zero. Therefore, there is a positive probability that excluded agents may never return to the non-excluded fold. The long-term distribution of types, in that case, is degenerate at $\mu^G = 0$.

Front-loading punishment is thus but one of countless efficient shapes for exclusion policies when amnesty is possible. When amnesty is not feasible, on the other hand, our efficiency criterion takes into account the welfare of existing agents who happen to be excluded at date 0. This restores benefits for front-load punishment, as the next result shows.

Proposition 9.1. *Let n^* be the highest number of periods for which borrowers alive at date 0 have been excluded, i.e. $n^* = \sup\left\{n : \mu_0^B(n) > 0\right\}$. Then, at any efficient policy, either $\phi_{n^*} = 1$ or $\phi_n = 0$ for all $n < n^*$.*

Proof. Among agents alive at date 0, the welfare of investors in good standing and investors who just lost their good standing ($n = 0$) is pinned down by $V^B(0)$. Therefore, these agents are indifferent across all exclusion policies that deliver $V^B(0)$. Assume that $\mu^B(1) > 0$ and consider any policy such that that $\phi_1 < 1$ while $\phi_0 > 0$. Recall that we must have

$$V^B(0) = \phi_0 V^G + (1 - \phi_0)V^B(1).$$

Holding all value functions the same then, we can lower ϕ_0 while maintaining $V^B(0)$ by raising $V^B(1)$, leaving the welfare of investors in good standing and those who just lost their standing ($n = 0$) unaffected, but making agents who have been in bad standing for one period ($n = 1$) strictly better off. In addition, since $V^B(0)$ is unchanged, any set of time-invariant equilibrium loan terms remain so, so that V^G is unaffected as well. This means that any policy such that $\phi_1 < 1$ while $\phi_0 > 0$ is suboptimal if $\mu^B(1) > 0$. This proves the result when

given an exclusion policy. The second step looks for policies that maximize V^N given $V^E(0)$. Our code is available for download at http://erwan.marginalq.com/index_files/wp.htm.

$n^* = 2$. Extending the argument recursively establishes the proposition for all n^*. □

The argument above begins with the same observation that leads to indeterminacy in the case with amnesty: under any SSE, the lifetime utility of agents who are in good standing at date 0 is fully summarized by $V^B(0)$. The same holds, obviously, for agents who just lost their standing ($n = 0$). Those agents are indifferent across all policies that deliver $V^B(0)$. Policies that front-load punishment ($\phi = 0$) on the other hand, benefit agents who have been excluded for more period ($n > 0$) since more punishment early means less necessary punishment later. This result has the following key implication for our purposes:

Corollary 9.2. *If $\mu_0^B(n) > 0$ for all $n \leq n^*$, then all efficient exclusion policies must have the shape described in proposition 4.1 up to n^*.*

Proof. If $\phi_{n^*} < 1$ at all optimal policies then the result follows. If $\phi_{n^*} = 1$ then repeating the argument above implies that either $\phi_{n^*-1} = 1$ or $\phi_n = 0$ for all $n < n^* - 1$, and the result follows. □

The premise that $\mu_0^B(n) > 0$ for all $n \leq n^*$ should generally be expected to hold since any policy that induces a stationary distribution such that $\mu^B(n) > 0$ for some n must also induce $\mu^B(n-1) > 0$. If initial conditions at date 0 are the result of an exclusion policy that has been in place for a while, they will therefore satisfy that premise.

In summary, front-loading punishment is only one of many efficient policies when amnesty is feasible. But some of front-loading punishment is once again strictly optimal when amnesty is not an option.

9.4 Heterogenous default costs and exclusion length

Assume that ex-post default costs can either be low at $\tau_L = \tau - \epsilon$ or high at $\tau_H = \tau + \epsilon$ where $\epsilon > 0$ and, for concreteness, assume that these two outcomes are equally likely. The exclusion length can be set so as to dissuade both ex-post types from defaulting from strategic reasons. Instead, it can be set to dissuade just the high-default cost borrower. Finally, exclusion length can be such that it does not dissuade either borrower type from defaulting for strategic reasons.

In other words, there are three possibilities. We can set κ to solve (6.4) for $\tau = \tau_H$ in which case only low-default cost borrowers default for strategic reasons. Low-default cost agents are then excluded for the corresponding time but since they cannot be dissuaded from strategic default, it makes no sense to exclude them any longer than what is strictly necessary to keep high-cost agents in line. If this option turns out to be optimal, note that imposing a mean-preserving spread on F results in lowering the length of exclusion. Second, we can set κ to solve (6.4) for $\tau = \tau_L$ so that no agent ever defaults for strategic reasons. In that case, the mean-preserving spread results in lengthening the duration of exclusion. Third and finally, we can simply give up on dissuading any agent from strategic default by setting $\kappa = 0$.

Each of those three results of spreading F is efficient for certain parameters. This means that, in general, mean-preserving spreads on incentives to default for strategic reasons have ambiguous effects on efficient exclusion length. We can describe this ambiguity more precisely.

Proposition 9.3. *Starting from an economy with homogenous default costs in which optimal exclusion length is positive, a mean-preserving spread in default costs raises exclusion length for ϵ small enough but must eventually drive exclusion length to zero as ϵ becomes large*

Proof. Start from the homogenous economy and introduce an infinitesimal spread $\tau_H - \tau_L = \epsilon > 0$. Adjusting κ by setting $\tau = \tau_L$ in (6.4) has no first order effect on any policy. Not adjusting, however, would cause half of agents with successful projects to begin defaulting for strategic reasons. Therefore adjusting by raising exclusion length infinitesimally is efficient. Once $\tau_H - \tau_L$ becomes large, high-default cost agents need not be dissuaded any longer while low default cost agents cannot be dissuaded by exclusion as τ_L becomes low and then eventually negative (these agents get positive utility from defaulting.) We now have no choice but to give up on the low-cost agents.[17] This completes the proof. □

Local mean-preserving spreads in default costs cause exclusion length to increase because it is efficient to keep low-default cost borrowers from defaulting for strategic reasons. But as the spread in F becomes large, exclusion threats become less potent. High default-cost agents do not default anyway while very low-default cost agents simply cannot be dissuaded from doing so.

[17] As the preceding discussion explained, before reaching zero there may be a point where it is optimal to only dissuade high-cost agents. Once that stage is reached, a bigger spread starts lowering exclusion length.