$u^{\scriptscriptstyle b}$

b UNIVERSITÄT BERN

Faculty of Business, Economics and Social Sciences

Department of Economics

Optimal Exclusion

Cyril Monnet Erwan Quintin

18-14

June 2018

DISCUSSION PAPERS

Optimal Exclusion

Cyril Monnet University of Bern and SZ Gerzensee Erwan Quintin^{*} Wisconsin School of Business

June 5, 2018

Abstract

In a canonical model of borrowing and lending, an exclusion technology that features full exclusion for a deterministic number of periods following default maximizes stationary equilibrium welfare. This exclusion policy maximizes the stationary volume of mutually beneficial lending transactions. It also maximizes the average welfare of the excluded. The optimal length of exclusion depends on fundamentals such as borrower patience and the direct cost of default. It also depends on incentives to default for strategic rather than exogenous reasons.

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 Schmidt-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

In this paper, we fully characterize the optimal shape of exclusion policies in a canonical model of lending with endogenous and exogenous default. A benevolent social planner who can design exclusion policies however she wishes chooses to exclude defaulting agents for a deterministic number of periods. While our social planner generally finds it optimal to inflict some punishment via exclusion in order to discourage strategic default, she also designs the exclusion policy so that excluded agents are allowed back into markets as fast as possible.⁴

This "harsh but short" punishment approach has two distinct advantages. First, front-

¹Two exceptions are Bond and Krishnamurthy (2004) and Corbae et al. (2016). Bond and Krishnamurthy (2004) model exclusion as direct constraints on transfers from lenders to investors and study the optimal shape of those constraints. However, in the context of our model 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.

 $^{^{2}}$ 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 interesting dynamics. Bethune et al. (2017) 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 noncooperative collusion with private information, Green and Porter (1984) show that short periods of non-cooperative plays are necessary to maintain cooperation.

⁴The corporate finance literature also shows that optimal financial contracts limit punishment or exclusion. For example, Bolton and Scharfstein (1990) show that lenders discipline borrowers by excluding them probabilistically from future loans. We add to this literature by showing that it is optimal to bring back excluded borrowers as fast as possible.

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. Because of those two complementary virtues, complete but finite exclusion is optimal no matter how the planner chooses to weigh the welfare of the excluded. Remarkably, maximizing the welfare of the excluded also maximizes the number of mutually beneficial transactions in stationary equilibrium.

Full but temporary exclusion following default, which we find to be optimal, is a good 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.⁵ 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 the utility they expect in autarky. One interpretation of this result is that 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 agents default in some cases because they have no choice, while

 $^{{}^{5}}$ 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.

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. 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, a social planner chooses to forgive defaulters in finite time. In fact, as we mentioned above, our benevolent planner chooses to design an exclusion policy that minimizes the expected 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, punishment is independent of history and, optimally, decreases in harshness through time. In that setting, lowering the severity of punishment 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, a planner has the ability to commit to any punishment profile she wishes and chooses to front-load punishment because any early forgiveness must be compensated for disproportionately in the future. Another key difference is that we consider a planner who must rely on markets in the spirit of Kehoe and Levine (1993) hence internalize the effect of her policy choices on contract terms.

Having fully characterized the optimal shape of the exclusion technology, we go on to show that the optimal length of exclusion depends on the model's fundamentals. Intuitively, the opportunity cost of exclusion depends on the profitability of investment projects, while the direct cost of exclusion depends for instance on agents' patience. We show that a higher average propensity to default for strategic reasons is associated with a higher optimal length of defaults. This makes intuitive sense since exclusion only serves to deter strategic default.

Importantly, this does not imply that the frequency of strategic default – an empirical

⁶Elul 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.

object about which there is an intense debate⁷ – is a good guide for choosing the optimal length of default. As has been pointed out in related contexts by Dubey et al. (2005) and Quintin (2013), there is no predicable relationship between the harshness of punishment and the frequency of strategic default. What we show is that the optimal shape of the exclusion technology can be characterized fully.

2 The environment

Consider an economy in which time is discrete and infinite. There is one good that cannot be stored across periods. The economy contains a mass one of infinitely-lived investors. Investors are each endowed with a project but no good. They can activate their project in each period by investing one unit of the good at the start of a period. When it is successful, with probability $\pi > 0$, a project pays output y > 0 at the end of the period and nothing otherwise. A law of large numbers holds so that π is also the fraction of projects that deliver positive output in a given period.⁸

To simplify the exposition, we will first assume that investors are risk-neutral and discount future payoffs at a constant period rate of $\beta \in (0, 1)$. However, linearity plays no role in any of our results. We relax it fully in section 6. There we consider an environment with risk-averse investors and traditional risk-sharing contracts. While this makes notation and derivations more burdensome, our results are unchanged.

Each period $t \in \{0, 1, ...\}$ a large mass of lenders are born endowed with a unit of the unique good. Lenders consume at the end of the period and then die. They can store the endowment they receive at the start of the period for a time-invariant safe payoff $R \in (0, y)$. They can also lend their endowment to an investor (or, equivalently, to a collection of lenders) in exchange for a promised payment $m_t \in [0, y]$ to be delivered if the project is successful at the end of date t. We will assume that lenders behave competitively in the sense that they take m_t as given.⁹ The fact that lenders live only for one period implies that all contracts

⁷See e.g. Foote et al. (2008) and Gerardi et al. (2015).

⁸In Bond and Krishnamurthy (2004), lenders find it optimal to exclude defaulting borrowers from future loans as long as they have not been reimbursed. In our context this essentially means permanent exclusion because investors must borrow in order to produce.

⁹While they could offer a contract with a higher payment no investor would take it. As is well known (see Quintin, 2012) incentives could exist in a model with endogenous default such as ours to offer a contract with a lower required payment but a lower probability of default. Below we will only consider equilibrium that make that option unprofitable for lenders. More simply, one can also apply the equilibrium concept of Dubey

must be one-period risky debt contracts which simplifies the analysis.¹⁰

Only investors observe their output realization. Furthermore, commitment on their part is limited. If they abstain from making a payment they experience a penalty τ drawn at the end of the period from a known distribution F. The default cost τ captures any remorse or shame defaulting investors experience as well as the effects of punishments other than exclusion. One can also think of \$ Absent any exclusion threats, investors of type $\tau \geq 0$ repay their loan when and only when their project is successful and

$$m_t \leq \tau.$$

Our environment captures in a simple way the two types of default events that correspond to the classification used in applied work.¹¹ Strategic default occurs when the project is successful but, nonetheless, investors choose to experience disutility τ rather than honor their debt. Non-strategic default occurs when the project fails and investors have no choice but to default. This clean dichotomy will enable us to obtain sharp comparative statics in section 5. Only investors know why they failed to pay so that other agents cannot distinguish defaults types from one another. 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.¹²

From the point of view of lenders, expected returns on loans to investors equal storage returns when:

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

Indeed, only the mass $1 - F(m_t)$ of successful investors whose τ is below m_t repay. As Quintin (2012, 2013) discusses at length there are potentially many solutions to this equation because of the endogeneity of default. To avoid the associated complications, assume for concreteness that when there are several interior solutions to (2.1), the lowest one prevails. This is the solution that makes investors as well off as possible in the set of payments lenders are willing to accept.¹³

et al. (2005) and assume away the issues raised by Quintin (2012, 2013) on anonymity grounds.

¹⁰We conjecture that our main results will not change much once that assumption is relaxed because all our arguments are developed conditional on agents being in default.

 $^{{}^{11}}$ See e.g. Foote et al. (2008.)

¹²We could assume that agents who default for exogenous reasons also experience disutility cost $\tau \ge 0$ or, for that matter, a different cost from strategic defaulters, without any effect on any part of the upcoming analysis.

¹³When F is continuous equation (2.1) does not display jumps so that if an m exists such that $\pi(1 - F(m))m > R$ then a solution exists by the intermediate value theorem. When F is not continuous, jumps

Since we have yet to introduce an exclusion technology there are no meaningful dynamics in this environment and all equilibria are stationary given the selection rule we specified above. In fact, this version of our economy 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, if no strictly positive solution to (2.1) exists lenders store their endowments and no project is activated. If, on the other hand, the set of solution to (2.1) is not empty, a stationary equilibrium simply is the lowest value m in that set together with the associated repayment policies. One of these two types of equilibrium must exist.

3 Economies with exogenous exclusion technologies

Assume now that an exclusion technology is available. Formally, an exclusion technology is a sequence $\{\phi_s\}_{s=0}^{+\infty}$ of *forgiveness* probabilities. When an agent fails to pay in a particular period, he is precluded from participating in lending markets – that is, he is kicked off from the economy – with probability $1 - \phi_0$ while with the complementary probability he is treated like any other agent. After this first draw, agents who have been excluded for n periods are allowed to return to the economy with probability ϕ_n .

We assume that forgiveness draws are independent across time so that, for instance, the probability that an agent is going to be excluded for exactly n periods 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 excluded agents will never be allowed to return to the loan market.

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 ($\phi_s = 0 \forall s$) – i.e. permanent exclusion – is the case

can be filled up by allocating different borrowers to different ends of each jump so that, once again, a break even payment can be found if and only if $\pi(1 - F(m))m > R$ for some m > 0.

studied in Kehoe and Levine (1993). Chatterjee et al. (2006) specify constant forgiveness rates, i.e. $\phi_s = \phi \forall s$. The one-time forgiveness lottery of Elul and Gottardi (2014) 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.¹⁴

Given the possibility of exclusion, investors can now be in different states at the start of a given period. Some agents are in default i.e. excluded, the rest have no default on their records and are allowed to participate in markets. Among excluded agents, we also need to record the number of periods for which they have been excluded since, in general, the forgiveness probability may depend on that length of time.

Throughout this paper we will focus on stationary equilibria, i.e. equilibria in which contract terms m and the mass of agents of each type are constant. In any such equilibrium, let V^N be the expected lifetime utility of investors who are not in default while $V^E(n)$ is the same for agents who have been excluded for n periods. Then, given the time-invariant payment m expected from investors,

$$V^{N} = (1 - \pi)\beta V^{E}(0) + \pi E_{\tau} \max\left\{y - m + \beta V^{N}, y - \tau + \beta V^{E}(0)\right\}$$
(3.1)

Clearly then, investors choose to pay once they discover their default cost τ if

$$m \le \tau + \beta \left[V^N - V^E(0) \right] \tag{3.2}$$

so that the lender's break-even condition becomes

$$\pi \left(1 - F\left(m - \beta \left[V^N - V^E(0)\right]\right)\right) m = R.$$
(3.3)

To complete the description of the environment we only need to define $V^{E}(n)$ for all n:

$$V^{E}(n) = \phi_{n} V^{N} + (1 - \phi_{n})\beta V^{E}(n+1).$$
(3.4)

Standard arguments show that, given m, equation (3.1) and equation (3.4) define a contraction

¹⁴We are making the implicit assumption that non-excluded investors are all treated equally once they default. This is without loss of generality as long as the stigma shock is independent across periods. But if types are persistent, welfare may be improved by making punishment depend on the credit history of investors.

mapping in V^N . This means that V^N is uniquely defined so that, in turn and by (3.1), $V^E(n)$ is uniquely defined for all n.

Several intuitively obvious features of these value functions will help the upcoming analysis. First, computing V^N only requires knowing $V^E(0)$ and m. In fact and by the same logic, given $V^E(0)$ there exists a lowest solution m to condition (3.2) which, among all solutions, makes the welfare of investors highest. Third, given $V^E(0)$, the optimal default policy hence the stationary rate of default can be computed. In the next section, these facts will enable us to split the search for the optimal exclusion policy into two natural steps. Finally, condition (3.4) implies that $V^E(n) \leq V^N$ for all n so that, holding $V^E(0)$ the same, welfare is highest when the mass of excluded agent is lowest.

To build a stationary equilibrium, start from m small enough that the left-hand side of equation (3.3) is below R (start from m = 0, say.) Then increase m until that left-hand side is R. If no $m \leq y$ meets that condition, the only stationary equilibrium is the no-trade equilibrium. In the other case, let $1 > \delta^D \geq 1 - \pi$ be the time invariant fraction of investors who default on their loan in a given period.

Let μ^N be the mass of investors who enter a period as non-excluded. In a stationary equilibrium, that object must satisfy:

$$\mu^{N} = \mu^{N} (1 - \delta^{D}) + \sum_{n=0}^{+\infty} \mu^{E}(n)\phi_{n}$$
(3.5)

where $\mu^{E}(0)$ is the mass of investors who defaulted in the previous period hence just entered exclusion while, for n > 0,

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

The following result tells us which exclusion technologies can support stationary equilibria with positive investment.

Lemma 3.1. A stationary equilibrium with $\mu^N > 0$ exists only if

$$\sum_{n=0}^{+\infty} \prod_{s=0}^{n} (1 - \phi_n) < +\infty.$$
(3.6)

Proof. Assume that $\sum_{n=0}^{+\infty} \prod_{s=0}^{n} (1-\phi_n) = +\infty$ but that $\mu^N > 0$. For all n,

$$\mu^{E}(n) \ge \mu^{E}(0) \prod_{s=0}^{n} (1 - \phi_{n}).$$

But

$$\mu^E(0) \ge \mu^N \delta^D$$

since the mass of newly excluded investors includes those who where not excluded at the start of the previous period and defaulted. (As we will discuss in the next section, $\mu^{E}(0)$ 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:

$$\sum_{n=0}^{+\infty} \mu^{E}(n) \ge \mu^{N} \delta^{D} \sum_{n=0}^{+\infty} \prod_{s=0}^{n} (1-\phi_{n}) = +\infty.$$

But in any stationary equilibrium we need

$$\mu^N + \sum_{n=0}^{+\infty} \mu^E(n) = 1$$

The result follows.

To understand this result, assume for instance that

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

among other ways in which we may have

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

Then a strictly positive fraction of investors who are excluded never return to markets. This implies that the only steady state of the Markov chain defined by the transition functions above puts all mass on the exclusion state. Hence, the only exclusion technologies that are compatible with strictly positive investment in equilibrium are the ones for which excluded investors eventually return to markets. Together with the condition we derived in the previous

 _		

section, that restriction also becomes sufficient to sustain positive investment in steady state.

Proposition 3.2. A stationary equilibrium with strictly positive investment exists if and only if:

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

2. A solution $m \leq y$ exists to equation (3.3).

In summary, given an exclusion technology so that excluded investors return to markets with probability one, an equilibrium only requires that a payment size be such that the expected return on loans equals the return on storage.

4 Optimal exclusion

Consider a benevolent planner who can select the exclusion technology before trade takes place. Above we showed that a stationary equilibrium exists for any given exclusion technology. Furthermore, as long as

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

a stationary equilibrium with strictly positive investment may exist. When there are multiple m that satisfy equation (3.3), the corresponding equilibria can be welfare ranked. The equilibrium with the lowest m maximizes V^N and $V^E(n)$ for all n. Since we take the point of view of a benevolent planner in this section, we assume that she has the ability to select the stationary equilibrium with the lowest m given the exclusion technology.

We will first consider a planner who weighs all investor types equally although many of the claims we make below generalize to cases where the planner values types differently, as we will discuss later. Assume then that the planner chooses an exclusion technology to maximize:

$$\mu^{N}V^{N} + \sum_{s=0}^{+\infty} \mu^{E}(s)V^{E}(s)$$
(4.1)

where, recall, μ^N is the mass of active investors in stationary equilibrium whereas, for all s, $\mu^E(s)$ is the mass of investors who have been in default for s periods. Equivalently, the social planner is maximizing the stationary equilibrium welfare of investors who will be randomly assigned to one of the countable types of investors.

The weights in the objective function sum up to one since the population of investors is one. Given this observation we can rewrite the objective function as

$$V^{N} + \sum_{s=0}^{+\infty} \mu^{E}(s) [V^{E}(s) - V^{N}].$$

This is intuitive: the welfare loss associated with excluding investors is $V^{E}(s) - V^{N}$ for each excluded type s.

This problem can be separated into natural blocks which allows us to solve this problem in several tractable steps. First, we study the optimal way to deliver a particular initial default value $V^{E}(0)$. Conditioning on $V^{E}(0)$ pins down the optimal m and V^{N} , by equations (3.1) and (3.3). The planner's goal now becomes:

$$\max_{\{\phi_s\}_{s=0}^{+\infty}} \sum_{s=0}^{+\infty} \mu^E(s) [V^E(s) - V^N]$$

subject to

$$V^{E}(0) = \phi_{0}V^{N} + (1 - \phi_{0})\phi_{1}\beta V^{N} + (1 - \phi_{0})(1 - \phi_{1})\phi_{2}\beta^{2}V^{N} + \dots$$
(4.2)

To understand the constraint, recall that excluded investors do not have any endowment and do not trade so that they only enjoy positive utility once they return to the non-excluded fold. The objective function in the above problem can be rewritten as:

$$\mu^{E}(0)\left\{\left[V^{E}(0)-V^{N}\right]+(1-\phi_{0})\left[V^{E}(1)-V^{N}\right]+(1-\phi_{0})(1-\phi_{1})\left[V^{E}(2)-V^{N}\right]+\ldots\right\}$$

since for all s > 0 we have

$$\mu^{E}(s) = \mu^{E}(0) \prod_{i=0}^{s-1} (1 - \phi_{i}).$$

To proceed consider the maximization problem that remains after dropping $\mu^{E}(0)$. In other words, consider the problem of maximizing the *average* welfare of excluded investors in stationary equilibrium. Using expression (3.4), the resulting objective is

$$\mathcal{P} = -(1-\phi_0) \left[V^N - \beta V^E(1) \right] - (1-\phi_0)(1-\phi_1) \left[V^N - \beta V^E(2) \right] - \dots$$

= $-V^N [(1-\phi_0) + (1-\phi_0)(1-\phi_1) + (1-\phi_0)(1-\phi_1)(1-\phi_2) + \dots]$
+ $(1-\phi_0)\beta V^E(1) + (1-\phi_0)(1-\phi_1)\beta V^E(2) + (1-\phi_0)(1-\phi_1)(1-\phi_2)\beta V^E(3) + \dots$

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

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

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 defaulting agents 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. The planner could even out forgiveness chances across excluded investors of each type. For instance, she could select the unique constant ϕ that meets the punishment constraint. This would give defaulting agents a chance to avoid exclusion altogether, for one potential benefit. But when the planner does so, she needs to increase punishment (reduce forgiveness odds) in the future. 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 $\mu^{E}(0)$ is the mass of agents that just became excluded at the start of a given period. In turn,

$$\mu^{E}(1) = (1 - \phi_0)\mu^{E}(0),$$

while

$$\mu^{E}(2) = (1 - \phi_0)(1 - \phi_1)\mu^{E}(0)$$

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

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

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 4.1. In any stationary equilibrium with positive investment,

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

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

Proof. In any stationary equilibrium,

$$\mu^{E}(0) = \left[\mu^{N} + \sum_{n=0}^{+\infty} \mu^{E}(n)\phi_{n}\right]\delta^{D}.$$
(4.3)

To understand this expression for $\mu^{E}(0)$, 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

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

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

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

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

$$\mu^{N} + \mu^{E}(0)(1+\zeta) = 1$$

or, given the lemma we just established,

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

In particular and quite remarkably, 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 and regardless of how the social planner wishes to weigh the welfare of the excluded versus that of the included, a policy of full but finite exclusion maximizes stationary equilibrium welfare. The following proposition – our main result – collects these results.

Proposition 4.2. In any stationary equilibrium that maximizes average welfare, the 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 $s > s^* + 1$.

In summary, bringing excluded investors back into markets as fast as possible given that a certain level of punishment needs to be imposed in equilibrium is optimal because doing so has two distinct virtues. First, it maximizes the stationary mass of active investors. Second and much less intuitively, it also maximizes the average welfare of excluded investors. The second feature stems from the fact that if the planner chooses to be lenient early and releases some of the excluded investors it must reestablish incentives by increasing punishment on investors that have been excluded for longer. 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. Obviously, 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.

5 Exclusion length

Having now fully characterized the shape of the socially optimal exclusion policy, we can discuss the one remaining defining feature of the optimal policy: the duration of exclusion. Obviously, the need for punishment hence the length of exclusion depends on every fundamental characteristic of the environment. To make this clear, we will consider a particular version of our economy where the optimal exclusion length can be solved for analytically and where sharp comparative statics can be obtained. We will then explain why, in general, the relationship between optimal exclusion length, fundamental parameters and the equilibrium quantity of default must be ambiguous.

5.1 Homogenous default cost

Consider a version of our economy where the distribution F of default costs only puts mass at one point τ . If $\tau \geq y$ then no agent ever chooses to default and the optimal exclusion length is zero. Indeed, exclusion only serves to deter strategic default. In fact, if $\tau \geq \frac{R}{\pi}$ then lenders can set $m = \frac{R}{\pi}$, which exposes them once again to zero strategic default and enables them to break even. In that case, once again, the optimal length of exclusion is zero.

When $\tau < \frac{R}{\pi}$, on the other hand, the case we will study for the remainder of this section, agents 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, the optimal exclusion policy must be the one that makes payment among investors with successful projects just optimal.

To make this precise, given an exclusion technology, define $\kappa \equiv \frac{V^E(0)}{V^N}$ as the fraction of active investor utility V^N borrowers give up when they choose to default. Given *and only given* the facts we have established about the shape of exclusion in this paper, κ is monotonically related to (= a sufficient statistic for) exclusion length. A lower κ means longer exclusion. We will refer to $\kappa \in [0, 1]$ as the exclusion discount. Borrowers pay when

$$m \le (1 - \kappa)\beta V^N + \tau. \tag{5.1}$$

It is clearly optimal for the planner to design an exclusion policy that makes this inequality tight. Otherwise she would be giving up on valuable trades by excluding investors for longer than is strictly needed. Furthermore, it is not just incentive compatible but also socially optimal for investors to pay since this makes $\mu^N = 1$. The break-even condition for lenders becomes:

$$\pi m = \pi ((1 - \kappa)\beta V^N + \tau) = R, \qquad (5.2)$$

and we can solve for V^N as follows:

$$V^{N} = \pi \left(y - m + \beta V^{N} \right) + (1 - \pi)\beta \kappa V^{N}$$

since, in this case, investors only default for non-strategic reasons.

This equation together with the facts that

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

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

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

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

$$\kappa = \frac{\pi^2 (y - \tau) - \frac{1}{\beta} (R - \tau \pi)}{\pi^2 (y - \tau) - (R - \tau \pi)}$$
(5.3)

The following proposition collects and summarizes the consequences of these results:

Proposition 5.1. When default costs are homogenous at given value τ , the optimal exclusion discount solves

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

In particular, exclusion length falls with investor patience (β) , project size (y), project quality (π) , and with the direct punishment (τ) associated with default.

An increase of patience increases the cost of exclusion and enables the planner 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 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.

This homogeneous version of our environment has the stark feature that no strategic default takes place in equilibrium. This is no longer the case when F is non-degenerate, a case to which we now turn.

5.2 Heterogenous default costs

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, we assume that these two outcomes are equally likely. This amounts to imposing a mean-preserving spread on the version of F we used in the previous subsection. The planner can choose to set the exclusion length to dissuade both ex-post types from defaulting from strategic reasons. Instead, she could dissuade just the high-default cost borrower, or, finally, she could choose to dissuade neither borrower type.

In other words, there are three possibilities. First, the planner can choose to set κ to solve (5.3) for $\tau = \tau_H$ in which case only low-default cost borrowers default for strategic reasons.

Low-default cost agent 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, the planner can choose to set κ to solve (5.3) 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, the planner can simply give up on dissuading any agent from strategic default by setting $\kappa = 0$.

It is easy to select parameters so that any of those three results of spreading F may solve the social planner's problem. This means that, in general, mean-preserving spreads on incentives to default for strategic reasons have ambiguous effects on optimal exclusion length. We can describe this ambiguity more precisely.

Proposition 5.2. 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 drives 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 (5.3) 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 optimal.

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.) The social planner now has no choice but to give up on the low-cost agents.¹⁵ This completes the proof.

Local mean-preserving spreads in default costs cause exclusion length to increase because the social planner finds it optimal 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.

¹⁵As the preceding discussion explained, before reaching zero there may be a point where the planner chooses to only dissuade high-cost agents. Once that stage is reached, a bigger spread starts lowering exclusion length.

6 Extensions

This section briefly considers five variations on our basic environment. The first three – riskaversion, exogenous punishment while excluded, and exogenous exit – have no effect on our key results about the shape of the optimal exclusion policy. The final two do change the nature of optimal exclusion in significant ways.

6.1 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 strictly 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. To allow for that possibility, write $m_H \in [0, y]$ for the payment from the investor to lender when the project succeeds while $m_L \geq 0$ is a payment from the lender to the investors when the project fails.¹⁶ As has been the case throughout this paper, we continue to focus on stationary equilibria. Investors now pay when the project is successful and given their default cost τ if

$$U(y - m_H) + \beta V^N \ge U(y + m_L - \tau) + \beta V^E(0)$$

where

$$V^{N} = (1-\pi) \left(U(m_{L}) + \beta V^{N}(0) \right) + \pi E_{\tau} \max \left\{ U(y-m_{H}) + \beta V^{N}, U(y+m_{L}-\tau) + \beta V^{E}(0) \right\},$$

while

$$V^{E}(0) = \phi_{0}V^{N} + (1 - \phi_{0})\left[U(0) + \phi_{1}\beta V^{N}\right] + (1 - \phi_{0})(1 - \phi_{1})\left[\beta U(0) + \phi_{2}\beta^{2}V^{N}\right] + \dots$$

¹⁶We could restrict contracts to those that stipulate $m_L = 0$ so that lenders do not provide any insurance at all to investors. Our point in this section is that the class of exclusion policy we have described in this paper is optimal regardless of whether or not risk-sharing happens in stationary equilibrium.

The expression for $V^E(0)$ is exactly the same as before except that it is potentially shifted up or down when $U(0) \neq 0$. As a result, it should be clear that the optimal shape of the exclusion policy cannot change. For one thing, a normalization of U(0) to zero is without loss of generality as long as $|U(0)| < +\infty$, which brings us back to the exact same expression as in the linear case.

A few more lines of algebra will clarify these facts and enable us to generalize our results to broader interpretations of U(0) in the next subsection. Write

$$V^{E}(0) = \frac{U(0)}{1-\beta} + \phi_{0}\widetilde{V}^{N} + (1-\phi_{0})\phi_{1}\beta\widetilde{V}^{N} + (1-\phi_{0})(1-\phi_{1})\phi_{2}\beta^{2}\widetilde{V}^{N} + \dots$$

where

$$\widetilde{V}^N = V^N - \frac{U(0)}{1 - \beta}.$$

A similar expression obtains for $V^E(n)$ for all n > 0. In other words, introducing general preferences results in level shift of exclusion 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^N . It follows that the same exclusion policy as in the linear case continues to maximize the average welfare of the excluded. The fact that this also maximizes the volume of transactions follows from the same arguments as before.

Algebraic considerations aside, the argument is unchanged because of the fundamental separation result this paper has established and built on. Once a equilibrium level $V^E(0)$ is set, it is optimal to administer the associated punishment as fast as possible, both because this makes the numbers of transactions as high as possible and because it maximizes the welfare of the excluded.

6.2 Exogenous punishment while excluded

Assume that the social planner is able to impose exogenous punishment¹⁷ on excluded investors which amounts to making the period utility excluded investors enjoy below U(0). Again, doing so amounts to shifting $V^E(0)$ up or down which cannot change the optimal shape of the exclusion policy. To see this, replace U(0) with some arbitrary, finite but possi-

¹⁷One practical and important form of punishment borrowers may experience while excluded from financial markets is wage garnishment. While comparatively rare in practice in the United States, garnishment is frequent and strictly enforced in most European economies. Livshits et al. (2007) study the quantitative importance of garnishment.

bly negative punishment utility U^P . The algebra we presented in the previous subsection is unaffected, and we once again reach the conclusion that the optimal shape of the exclusion policy is unaffected.

A more complex problem results from assuming that U^P can change over time. With enough freedom to select the level of exogenous punishment, one can justify exclusion policies that differ drastically from the punish-as-fast-as-possible solution. For instance, it becomes possible to threaten defaulters with very large punishment with low probability in the future. Positive levels of early forgiveness may then become optimal, compensated by a non-zero probability of reaching a distant but high exogenous punishment period.

6.3 Exogenous exit

Assume that that an exogenous fraction χ of excluded investors exogenously return to the non-excluded set of agents. This is equivalent to assuming that a fixed fraction of agents die each period and are replaced by an offspring whose welfare they value exactly as they do their own, with, in particular, the same discount rate. Given exogenous exit from exclusion, the forgiveness policy must feature $\phi_n \geq \chi$ for all n which directly implies that condition (3.6) holds. The exact same argument as in the unconstrained case implies that $\phi_n \geq \chi$ is binding until the desired level of punishment has been reached.

In other words, the substance of our result – that punishment should be front-loaded as much as feasible – does not change when exogenous exit is introduced. The case where exiting agents do not value the welfare of their replacement like their own is more tedious to deal with because this introduces a wedge between the social planner's valuation of exit and that of investors. But it remains the case that, optimally, punishment via exclusion should be front-loaded.

6.4 Observable income

Consider a version of our model in which project outcomes are public information. This allows the social planner to recognize whether default occurred by choice or by constraint. She can now condition exclusion length on default types. A first, obvious consequence of this change is that the social planner will immediately forgive all exogenous default. Indeed, that type of default cannot be helped and excluding those agents would result in needlessly lowering the volume of lending contracts in stationary equilibrium.¹⁸

On the other hand, the social planner will usually choose to exclude agents who default for strategic reasons, especially since she can now do so without wasting the associated incentives on borrowers who do not need them. If a level of exclusion exists that completely eliminates strategic default, imposing that level is obviously optimal. At the other extreme, it is possible (see subsection 5.2 for specific examples) that strategic default cannot be dissuaded in which case exclusion is once again socially wasteful. In general however, our model clearly predicts that when conditioning is feasible, borrowers who can document extant exogenous shocks should be excluded for a shorter period.

6.5 Non-stationary policies

Assume that our planner, instead of maximizing long-term stationary equilibrium welfare, maximizes the average welfare of the current investors, with type arbitrarily distributed at date 0. Assume that they have the ability to commit to an entire future path of exclusion policies. That is, they have the ability to announce and commit to forgiveness probabilities for every excluded type at every future date.

One immediate consequence of this change in objective is that a stationary solution is no longer optimal. First and most obviously, the planner will always choose to forgive all the excluded investors at date 0 since punishing them has no impact on incentives while forgiving them makes the volume of transactions as high as it can be at date 0. After date 0 however, the incentives we have discussed throughout this paper are active and the planner will generally find it optimal to commit to some exclusion following default in the future.

If the planner does not have the ability to commit, by the same logic, the set of feasible/sustainable exclusion policies shrinks even more drastically. Indeed, forgiving currently excluded agents becomes optimal at every history. Although making this point formally requires solving an intricate repeated game, a natural conjecture is that the only subgame perfect outcome in the resulting environment is full forgiveness in every period which brings us back to the environment we described in section 2 where only direct, exogenous punishment

¹⁸This feature of our environment has a practical counterpart in the fact that the underwriting process for new mortgages often inquires about and documents causes for recent blemishes. Loan officers and mortgage underwriters are more likely to extend new loans to impaired borrowers when they can establish that exogenous circumstances caused the credit difficulties on record, provided that those exogenous circumstances have been resolved, as is the case by assumption in our model. Persistent credit impairment shocks do tend to be associated with lengthy exclusion and this would happen de facto in our model as well.

can sustain lending.

7 Conclusion

In a canonical model of borrowing and lending with endogenous default, we find that an exclusion technology that features full exclusion for a finite number of periods maximizes stationary equilibrium welfare. Not only does this shape for the exclusion technology 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.

While we find that this result is robust to a host of considerations, we have abstracted from two important complications. As in Kehoe and Levine (1993) or Bulow and Rogoff (1989), our agents do not accumulate assets over time. This means that we do not need to keep track of an endogenous distribution of wealth as a state variable. A tractable extension with durable assets along the lines of Lagos and Wright (2005) would not change our result since the wealth distribution is degenerate in that case. Dealing with non-degenerate wealth distributions in more general environments is a promising if complex avenue for future work.

Likewise, we have focused on stationary equilibrium and exclusion policies. In so doing, not only do we reduce the dimension of our problem to a manageable size, we also sidestep the time-consistency issues that would arise if we allowed for the possibility of one-time amnesties, among other non-stationary options. Implicitly then we have focused on a social planner that has the ability to commit and cares about long-term welfare. To us, this seems a natural guiding principle for setting up the legal institutions that frame lending. Studying the consequences of limited commitment on the part of policy makers should also lead to interesting insights.

8 Appendix

8.1 Simplification of expression \mathcal{P}

We need to show that

$$(1 - \phi_0)V^E(1) + (1 - \phi_0)(1 - \phi_1)V^E(2) + \dots$$

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

$$\begin{aligned} &(1-\phi_0)\phi_1V^N &+(1-\phi_0)(1-\phi_1)\phi_2\beta V^N &+(1-\phi_0)(1-\phi_1)(1-\phi_2)\phi_3\beta^2 V^N + \dots \\ &(1-\phi_0)(1-\phi_1)\phi_2 V^N &+(1-\phi_0)(1-\phi_1)(1-\phi_2)\phi_3\beta V^N &+(1-\phi_0)(1-\phi_1)(1-\phi_2)(1-\phi_3)\phi_4\beta^2 V^N + \dots \\ &(1-\phi_0)(1-\phi_1)(1-\phi_2)\phi_3 V^N &+(1-\phi_0)(1-\phi_1)(1-\phi_2)(1-\phi_3)\phi_4\beta V^N &+(1-\phi_0)(1-\phi_1)(1-\phi_2)(1-\phi_3)(1-\phi_4)\phi_5\beta^2 V^N \\ &+\dots \end{aligned}$$

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

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

In order to support some investment in equilibrium, the planner has to restrict herself to exclusion policies such that the second term is zero.¹⁹ As a result, the first column gives $(1-\phi_0)V^N$. 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:

$$(1-\phi_0)V^N + (1-\phi_0)(1-\phi_1)\beta V^N + (1-\phi_0)(1-\phi_1)(1-\phi_2)\beta^2 V^N + \dots$$

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

$$V^{N} + (1-\phi_{0})\beta V^{N} + (1-\phi_{0})(1-\phi_{1})\beta^{2}V^{N} + \dots - \phi_{0}V^{N} - (1-\phi_{0})\phi_{1}\beta V^{N} - (1-\phi_{0})(1-\phi_{1})\phi_{2}\beta^{2}V^{N} - \dots$$

The second part of the expression, given the constraint, is simply $V^E(0)$. The first part, other than for the very first term, yet again features a common factor $(1 - \phi_0)$ which can be used to split it into two subparts, leaving us with:

$$V^{N} - V^{E}(0) + \beta V^{N} + (1 - \phi_{0})\beta^{2}V^{N} + \dots$$
$$-\beta \left\{ \phi_{0}V^{N} - (1 - \phi_{0})\phi_{1}\beta V^{N} - (1 - \phi_{0})(1 - \phi_{1})\phi_{2}\beta^{2}V^{N} - \dots \right\}$$

¹⁹In 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.

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

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

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

8.2 Full but finite exclusion maximizes \mathcal{P}

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

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

subject to the restriction that condition 4.2 must hold i.e. that $V^E(0)$ 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 4.2.

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:

$$\phi_0 + (1 - \phi_0)\phi_1\beta + (1 - \phi_0)(1 - \phi_1)\beta^2 = \frac{V^L(0)}{V^N}$$

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

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

The proposition holds if $\phi_0 > 0 \implies \phi_1 = 1$. 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:

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

Rewrite the constraint as

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

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

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

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 4.2 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

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

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

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

accounted for by $\{\bar{\phi}_t\}_{t=0}^{T-1}$ has to be within $\frac{\epsilon}{k}$ of $V^E(0)$. Now consider the alternative policy $\{\hat{\phi}_t\}_{t=0}^{T}$ 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 $V^E(0)$ 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 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 (2017), "Indeterminacy in Credit Economies," University of Virginia Working Paper.

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.

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.

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.

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.

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.

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," *Journal of Mathematical Economics*, 49, pp. 418–421.

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