CAEPR Working
Paper
#2014-002

Optimal Bankruptcy Code: A Fresh Start for Some

Grey Gordon
Indiana University

September 8, 2014

This paper can be downloaded without charge from the Social Science Research Network electronic library at <http://ssrn.com/abstract=2494888>

The Center for Applied Economics and Policy Research resides in the Department of Economics at Indiana University Bloomington. CAEPR can be found on the Internet at: <http://www.indiana.edu/~caepr>. CAEPR can be reached via email at caepr@indiana.edu or via phone at 812-855-4050.

©2014 by Grey Gordon. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Optimal Bankruptcy Code: A Fresh Start for Some

Grey Gordon*
Indiana University

September 8, 2014

Abstract

What is the optimal consumer bankruptcy law? I examine this question in the context of an incomplete markets lifecycle model with a planner who can choose state-contingent bankruptcy costs. I develop two key theoretical characterizations. First, the optimal policy has a bang-bang property: The planner either gives a household a “fresh start” or forbids it from filing. Second, it is optimal for the planner to always allow bankruptcy if the household cannot repay or would prefer an outside option. Consequently, a natural borrowing limit economy—an economy where bankruptcy is never allowed—is suboptimal. Quantitatively, the optimal policy results in large amounts of debt and default with ex-ante welfare gains, relative to a no-borrowing economy, as large as 12.8% of lifetime consumption. While the optimal policy is complicated, a simple cutoff rule allowing bankruptcy when a household’s debt is 2.6 times its endowment results in a welfare gain of 12.2%.

1 Introduction

Bankruptcy policy varies greatly by time and location. In many European countries, there is little to no debt forgiveness. Bankruptcy laws in the United States, on

*I thank Kartik Athreya, Satyajit Chatterjee, Daphne Chen, Bulent Guler, Aaron Hedlund, Juan Carlos Hatchondo, Dirk Krueger, Amanda Michaud, Victor Ríos-Rull, and Eric Young. I also thank conference participants from the European Meeting of the Econometric Society 2014, Computing in Economics and Finance 2014, Midwest Macro 2013 and 2014, and a UIUC mini-conference in 2014, as well as seminar participants at Florida State University and Indiana University.

the other hand, are widely considered pro-debtor. Moreover, views on the “proper” amount of debt forgiveness have changed dramatically over the last two hundred years. In the U.S., debtors’ prisons have been replaced with a relatively swift bankruptcy process, which, until recently, offered a near-complete discharge to almost everyone. In 2005, the Bankruptcy Abuse Prevention and Consumer Protection Act restricted this near-complete discharge to only those with below-median income, forcing above-median income households to pay all their disposable income (in the specific sense defined by the law) for five years.¹ Which of the many possible bankruptcy laws, ranging from complete discharge for all to no discharge at all, is best?

To answer this question, I use an incomplete markets lifecycle model of bankruptcy (Livshits, MacGee, and Tertilt, 2007) but allow a planner to choose state-contingent bankruptcy penalties. The planner specifies whether a household may file for bankruptcy, any pecuniary costs associated with filing, and how long a bankruptcy remains on a household’s credit record. In choosing these penalties, the planner faces a tradeoff between providing households with *intertemporal* consumption smoothing and *intra*temporal smoothing: By treating defaulters harshly, he improves credit markets and thereby reduces the cost of transferring resources across time; by treating bankrupts leniently, he adds state-contingency into debt repayment, allowing households to transfer resources across contemporaneous states.

I analytically characterize the planner’s optimal policy for a broad class of utility functions, social welfare functions, and endowments, holding the risk-free interest rate fixed. I find two main results. First, the optimal policy has a bang-bang property: The planner either gives a household a “fresh start” or forbids it from filing.² The fresh start in the optimal policy deviates substantially from current U.S. policy in that there are no filing costs and a bankruptcy filing never impacts a household’s credit record. Second, I find that a natural borrowing limit economy—an economy where households have maximal commitment to repay their debt—is inferior to one that allows bankruptcy when a household either cannot repay or would prefer an

¹Robe, Steiger, and Michel (2006) document both the historical and geographical variation in bankruptcy laws, and Coleman (1974) documents the variation in U.S. bankruptcy laws over time. As Wang and White (2000) summarize, “the United States is extremely unusual in having very prodebtor bankruptcy laws” (p. 255). For details on the Bankruptcy Abuse Prevention and Consumer Protection Act, see White (2007).

²The policy is not stochastic. Rather, it allows bankruptcy for households in some regions of the state space and forbids it in others. Households may stochastically transition into or out of the regions, but the regions themselves are fixed and common knowledge.

outside option. This outside option represents the best of any default options apart from bankruptcy, and can be set low enough to be irrelevant (as if not available to households).³

Using the analytic characterizations, I then examine the optimal policy quantitatively, focusing on parameter values calibrated to the U.S. economy. Relative to the worst possible outcome, which in the model environment is a zero borrowing limit economy, the optimal policy produces a welfare gain of 12.8% measured in consumption equivalent variation for newborn households. To put this gain in perspective, it is equivalent to increasing U.S. personal consumption expenditures by roughly 1.5 trillion dollars.⁴ In contrast, a bankruptcy law similar to that of the current U.S. system produces a welfare gain of 1.3%. Relative to the U.S., the optimal policy results in a four-fold increase in the bankruptcy filing rate and a thirty-fold increase in debt.

While very complicated, the optimal policy has three distinguishable features. First, bankruptcy is typically not allowed for the “luckiest” households (those with an endowment two standard deviations above the mean for their age). This feature lowers interest rates by effectively forcing lucky households to repay their debts. Second, bankruptcy is typically allowed for the largest debt amounts, but frequently not for moderate debt amounts. This prevents any household from having to make large debt payments via a drastic cut in consumption. It also keeps interest rates for moderate debt levels at around 15% to 36% annually. Third, the policy is non-monotonic, allowing bankruptcy for one state and not its neighboring state seemingly at random. In fact, the planner’s choices are not random at all, but are designed to allow households to accumulate significant debts while regularly discharging them as they transition through the state space.

The optimal policy, by definition, offers the best consumption smoothing to households, but its non-monotone aspect requires that households and creditors know intricate details of bankruptcy law and respond optimally. In practice, these informational requirements would make it an unmitigated disaster. This raises the question, “Is there a simpler policy that can come close to the optimal one?” As it turns out, there is: A rule allowing bankruptcy if and only if a household’s debt-endowment

³This irrelevance is formalized in Proposition 2, which shows households never choose the outside option in equilibrium if the value from it is set low enough.

⁴Personal consumption expenditures in 2013 amounted to 11.5 trillion dollars. Multiplying by 12.8% gives 1.5 trillion dollars.

ratio exceeds 2.6 results in a welfare gain of 12.2%. This rule comes close to the welfare gain of the optimal policy but is far simpler to understand.

The numbers reported thus far assume that the outside option is set low enough to be irrelevant. In the opposite extreme, I prove that when the outside option equals the value of autarky, the planner can do no better than the worst case of zero debt supported in equilibrium. Quantitatively, a plausible intermediate case results in the optimal policy producing a welfare gain of 5.3%, substantially less than the 12.8% produced when the outside option is very low. Moreover, a bankruptcy rule allowing bankruptcy if and only if the outside option would otherwise be preferable does as well as the optimal rule. These results suggest, quite intuitively, that to maximize welfare gains from bankruptcy law, one must limit households from defaulting outside the bankruptcy system.

1.1 Related Literature

The papers most closely related to this one are [Grochulski \(2010\)](#) and [Chen and Corbae \(2011\)](#).⁵ In the former, a social planner faces the Mirrleesian problem of delivering consumption subject to an incentive compatibility constraint (derived from truthful reporting of household income). The optimal policy is solved for and then implemented using bankruptcy. The focus in this paper is on full information and what the planner can legislate. The full information assumption is reasonable given that current U.S. bankruptcy law requires information disclosure. [Grochulski \(2010\)](#) also requires a two-state income process and a smooth utility function whereas the present paper allows any discrete Markov income process and any continuous, increasing utility function that is unbounded below. [Chen and Corbae \(2011\)](#) use a calibrated model to investigate the optimal length of time for a bankruptcy to remain on a household's credit record. Interestingly, they find that welfare monotonically decreases in the duration. This result is analogous to the present paper's theoretical finding that a duration of zero is optimal.

A large literature has examined various consumer bankruptcy policies, but, with the exception of [Grochulski \(2010\)](#), none has considered such a large menu.⁶ A

⁵[Wang and White \(2000\)](#) has similar objectives but a different approach featuring a two-period model. The optimal policy in [Wang and White \(2000\)](#) is qualitatively similar to the debt-endowment cutoff rule considered in Section 4.4.

⁶There is also a separate literature examining optimal firm bankruptcy law.

non-exhaustive list of examples includes [Li and Sarte \(2006\)](#); [Chatterjee, Corbae, Nakajima, and Ríos-Rull \(2007\)](#); [Livshits et al. \(2007\)](#); [Athreya, Tam, and Young \(2009a,b\)](#); [Chatterjee and Gordon \(2012\)](#); and [Gordon \(2013\)](#). The overriding theme in these papers is that harsher bankruptcy laws almost always improve welfare, even to the point that eliminating a bankruptcy option greatly improves welfare in the absence of shocks to household expenditures.

Beginning with [Li and Sarte \(2006\)](#), a number of papers have examined the effects of having a default option in addition to Chapter 7 bankruptcy.⁷ They have found that the inclusion of alternate default options substantially alters equilibrium behavior. The present paper’s inclusion of an outside option is meant to represent, in a reduced form way, a household’s best default option apart from Chapter 7 bankruptcy. By allowing the outside option value to be arbitrarily low or as large as the value of autarky, and by not placing any regularity conditions (such as monotonicity) on it, the option can capture many different types of default options. The theory assumes creditors receive nothing in the case of default, which does limit its generality. However, recovery rates on defaulted debt tend to be quite low: [Chatterjee and Gordon \(2012\)](#) report the average gross recovery rate on defaulted debt is around 20% while the net recovery rate is 12-14%.

[Dávila, Hong, Krusell, and Ríos-Rull \(2012\)](#) have recently presented a notion of optimality in incomplete markets where a planner chooses savings decisions. The equivalent here would be for the planner to choose both savings and default decisions. While interesting, this paper asks a more decentralized question: What is the optimal bankruptcy structure given the best response of households?

2 Model

The model is very similar to [Livshits et al. \(2007\)](#) with the main difference being state-contingent default costs.

⁷Examples dealing with consumer credit include [Chatterjee and Gordon \(2012\)](#), [Athreya, Sanchez, Tam, and Young \(2012a, 2014\)](#), and [Chen \(2013\)](#). There is also a literature looking at bankruptcy and foreclosure, with examples including [Li and White \(2009\)](#), [Li, White, and Zhu \(2011\)](#), and [Mitman \(2011\)](#).

2.1 Model Setup

The economy is populated with a continuum of households who live for at most T periods. The household labor endowment e is strictly positive and follows a finite-state Markov chain $\pi_{ee'|t}$ that is age-dependent.⁸ Newborn households draw their e from a distribution π_e .

Markets are incomplete with households only having access to a bond $a \in A$, with $a < 0$ being unsecured debt and $a \geq 0$ being savings. A is a finite set containing negative and positive elements, as well as zero. For simplicity, I assume households may not default on savings $a \geq 0$. Newborn households have $a = 0$.

Households have a bankruptcy flag h which indicates whether a household has a bankruptcy on record, $h = 1$, or not, $h = 0$. All households begin life with $h = 0$. It is assumed that when a household defaults or has a bankruptcy record that they may not borrow, $a' \geq 0$. This exogenous restriction is for tractability, and means only *one* bankruptcy law (since households will not default on $a \geq 0$) is needed rather than two.

Preferences are additively time-separable over consumption

$$\sum_{t=1}^T \beta^{t-1} u(c_t) \tag{1}$$

with a discount factor $\beta > 0$. The period utility u is continuous and strictly increasing, but not necessarily concave or differentiable. Additionally, it is unbounded below with $\lim_{c \downarrow 0} u(c) = -\infty$. Having utility unbounded below is important for the theoretical results. However, it does mean $u(0)$ is not defined, and so I require households choose $c > 0$.⁹ Note that the commonly used constant relative risk aversion (CRRA) preferences are also not defined at zero when CRRA is greater than one.

A household's state is (a, e, t, h) . The price of a discount bond with face value a' is $q_t(a', e)$. For savings, $a' \geq 0$, default is not allowed, and so the price for such a bond is simply equal to the risk-free rate of transferring resources across time, $\bar{q}$. For debt, $a' < 0$ (which implies $h = 0$), creditors expect a repayment rate $p_t(a', e)$.

⁸Other sources of heterogeneity, such as preference shocks or mortality risk, could be easily added. For the sake of exposition, this has not been done.

⁹Alternatively, a consumption minimum $\underline{c} > 0$ could be imposed with a requirement that $c \geq \underline{c}$. The choice set is compact either way because a' is restricted to lie in a finite set.

Consequently, a no arbitrage condition¹⁰ has

$$q_t(a' < 0, e) = \bar{q}p_t(a', e). \quad (2)$$

When the repayment rate p_t is low, so is the discount q_t implying a high interest rate $1/q_t - 1$. Equilibrium requires the repayment rate $p_t(a', e)$ is consistent with household default decisions. For convenience, I define $p_t(a' \geq 0, e) = 1$ so that $q_t(a', e) = \bar{q}p_t(a', e)$ for all a', e .

Households always have access to an informal default option, which I call an outside option, that delivers $V_t^O(a, e)$ measured in lifetime discounted utility terms. I make two key assumptions regarding the outside option. First, if a household exercises its outside option, the creditor gets nothing. Second, $V_t^O(a, e)$ is less than or equal to the value of autarky $V_t^A(a, e)$ for all a, e, t . I let the value of autarky allow for savings at the risk free-rate, and so it is given by $V_t^A(a, e) = X_t(0, e)$ where

$$\begin{aligned} X_t(a, e) &= \max_{a' \in A} u(c) + \beta \sum_{e'} \pi_{ee'|t} X_{t+1}(a', e') \\ \text{s.t. } &c + \bar{q}a' = e + a, \quad c > 0, a' \geq 0 \end{aligned} \quad (3)$$

with $X_T := u(e + a)$. The outside option represents the best of all a household's options to informally default. Proposition 2 shows the outside option can be set low enough so as to not affect households decisions (i.e., households behave as if they did not have access to an outside option).

Bankruptcy policy is defined by the instruments available to the planner. Specifically, the planner can do all of the following:

1. Specify whether a household is allowed to file for bankruptcy, $D_t(a, e) = \{0, 1\}$, or not, $D_t(a, e) = \{0\}$.
2. Confiscate any fraction $\chi_t(a, e) \in [0, \bar{\chi}]$ of a household's endowment in the period of filing with $\bar{\chi} < 1$.
3. Charge a lump sum bankruptcy filing cost $\zeta_t(a, e) \geq 0$.
4. Retain a bad credit record with probability $\lambda_t \in [0, 1]$.¹¹

¹⁰See Gordon (2013) for the interpretation of this common debt-pricing equation as a no-arbitrage condition.

¹¹It should be possible to relax this and allow λ_t to also be conditioned on a and e . However, this

Note the planner may not condition on the complete history of a household. While it rules out potentially interesting policies, such as those conditioned on the number of times a household has filed, it allows for many other interesting ones. The preceding abilities of the planner cover, as special cases, many of the types of punishment that the literature has used to model bankruptcy. Given the logic of the proofs, it would be easy to add additional policy instruments inflicting costs on the household budget constraint while in bankruptcy (not just in the period of filing); however, to keep the notation simple, this is not done.

If the household is allowed to file for bankruptcy and does so, then the planner must forgive all debt in exchange for all assets. That is, the planner cannot forgive only a fraction of the debt and make the household responsible for the residual. Relaxing this assumption would complicate the analysis because of its effect on debt pricing. It would also make the debt partially secured, but the focus of this paper is on unsecured credit.

2.2 Household Problem

Given the policy instruments and prices, the value function of a household with a, e, t who has no record of a bankruptcy is¹²

$$\begin{aligned} V_t(a, e, h = 0) &= \max_{o, d} oV_t^O(a, e) + dV_t^D(a, e) + (1 - o)(1 - d)V_t^R(a, e) \\ \text{s.t. } &o \in \{0, 1\}, d \in D_t(a, e), od = 0 \end{aligned} \quad (4)$$

where the value of repaying debt is

$$\begin{aligned} V_t^R(a, e) &= \max_{a' \in A} u(c) + \beta \sum_{e'} \pi_{ee'|t} V(a', e', 0) \\ \text{s.t. } &c + q_t(a', e)a' = e + a, \quad c > 0 \end{aligned} \quad (5)$$

is a bit unnatural as λ affects the household not only in the period of filing, but also in bankruptcy. One could add state variables to keep track of a household's state at the time of filing, but this would increase notation significantly without adding any insight.

¹²To cutdown on notation, the definition below treats bankruptcy and repayment both as feasible options. If bankruptcy is not feasible, then the value function should be reformulated without reference to V^D and similarly for V^R . Additionally, default options are only available if $a < 0$.

and the value of defaulting on debt via bankruptcy is

$$\begin{aligned}
V_t^D(a, e) &= \max_{a' \in A} u(c) + \beta \sum_{e'} \pi_{ee'|t} ((1 - \lambda_t) V_t(a', e', 0) + \lambda_t V_t(a', e', 1)) \\
\text{s.t. } &c + \bar{q}a' + \zeta_t(a, e) = e(1 - \chi_t(a, e)), \quad c > 0, a' \geq 0.
\end{aligned} \tag{6}$$

For $a \geq 0$, households are not allowed to file for bankruptcy.

If a household's budget constraint would be empty upon choosing to repay, they must choose either the outside option or bankruptcy (if an option). If bankruptcy is not a feasible option (say because ζ_t is too large), then they must choose the outside option or repay. Further, to handle the case $t = T$, define $V_{T+1} = 0$, and $q_T = 0$.

The value of having a bankruptcy record, $h = 1$, is similar to the value of default:

$$\begin{aligned}
V_t(a, e, h = 1) &= \max_{a' \in A} u(c) + \beta \sum_{e'} \pi_{ee'|t} ((1 - \lambda_t) V_t(a', e', 0) + \lambda_t V_t(a', e', 1)) \\
&c + \bar{q}a' = e + a, \quad c > 0, a' \geq 0.
\end{aligned} \tag{7}$$

Note that I do not allow households in bad standing access to an outside option. This is without loss of generality (wlog): $V_t^O(a, e)$ is less than the value of autarky and $V(a, e, h = 1)$ is greater than it.

2.3 Equilibrium

Steady state equilibrium for a given set of policy instruments, D_t , λ_t , χ_t , ζ_t and a risk-free price $\bar{q}$ is a set of policies, c_t, a'_t, d_t, o_t , value functions V_t , and prices q_t such that

1. Households optimize taking prices as given.
2. The price schedule q_t is given by no arbitrage:

$$q_t(a', e) = \bar{q}p_t(a', e). \tag{8}$$

3. Repayment rates p_t are consistent:

$$p_t(a', e) = \sum_{e'} \pi_{ee'|t} (1 - \max\{d_{t+1}(a', e'), o_{t+1}(a', e')\}). \tag{9}$$

4. The distribution of households is invariant.

2.4 Planner's Problem

The planner solves

$$\begin{aligned} & \max_{\chi, \zeta, \lambda, D} \sum_{a, e, t} \alpha_t(a, e) V_t(a, e, h = 0) \\ \text{s.t. } & D_t(a, e) \in \{\{0, 1\}, \{0\}\}, \lambda_t \in [0, 1], \chi_t(a, e) \in [0, \bar{\chi}], \zeta_t(a, e) \geq 0 \forall a, e, t \end{aligned} \quad (10)$$

where $\alpha_t(a, e) \geq 0$ is the weight placed on type $(a, e, h = 0, t)$. Note that I restrict the weights to be zero for $h = 1$. This is not to say the planner does not care about utility in the $h = 1$ states: He implicitly does if they are potentially visited by household having a state $(a, e, t, h = 0)$ with $\alpha_t(a, e) > 0$. However, one of the key proofs shows that for any policy with $\lambda_t > 0$ there is another policy with $\lambda_t = 0$ generating the same planner utility. Since the case of $\lambda_t = 0$ results in $h = 1$ being off the equilibrium path, restricting α_t to only place weight on $h = 0$ is convenient.

3 Theoretical Results

This section characterizes the optimal policy theoretically. To simplify the proofs, a household is assumed to repay when indifferent between bankruptcy and repayment. If indifferent between bankruptcy and the outside option, they are assumed to file for bankruptcy if it is an option. All proofs are relegated to Appendix A.

Proposition 1 establishes that an equilibrium exists for any planner policy. The proof is far simpler than in Chatterjee et al. (2007) because of the partial equilibrium assumption. It is also simpler than in Livshits, MacGee, and Tertilt (2003) because the household choice set is finite.

Proposition 1. *For any policy choice of the planner, an equilibrium exists.*

Lemmas 1 and 2 establish some basic properties of the household value function. Lemma 1 shows higher prices, i.e. lower interest rates, for debt improve household welfare. Lemma 2 shows households are weakly better off not having a bankruptcy record. The proof is straightforward since households have an option value of borrowing if they do not have a bankruptcy record.

Lemma 1. $V_t^R(a, e)$ is weakly increasing in $q_t(a', e)$ for any $a' \leq 0, e$ and all a .

Lemma 2. $V_t(a, e, h = 0) \geq V_t(a, e, h = 1)$ for all $a \geq 0$.

Lemma 3 shows the value function is bounded above and below when households have positive assets. It also shows these bounds are independent of planner policy and the outside option. This lemma is essential to the proof of Proposition 2.

Lemma 3. For $a \geq 0$, $V_t(a, e, h)$ is uniformly bounded above and below. Moreover, defining the maximum and minimum values possible for e as $\bar{e}$ and $\underline{e}$, these bounds depend only on $\underline{e}$, $\bar{e}$, $\min A$, $\bar{q}$, β , and u .

Given that most of the literature does not allow an outside option to be chosen, one may wonder how important allowing an outside option is. Proposition 2 shows it need not be important: If V_t^O is small enough, it will not be chosen in equilibrium.¹³ The proof involves showing households would not willingly subject themselves to the outside option with positive probability because doing so would violate the lower bound found in Lemma 3.

Proposition 2. There exists a δ such that $V_t^O(a, e) < \delta$ for all a, e, t implies the outside option is never chosen in equilibrium.

Lemma 4 establishes the continuity of V^D in the bankruptcy filing cost ζ . While intuitively obvious, the proof is non-trivial and uses the Inada condition on u .

Lemma 4. $V_t^D(a, e)$ is continuous in $\zeta_t(a, e)$ for any $\zeta_t(a, e) \in [0, e(1 - \chi_t(a, e))]$. Since bankruptcy is a feasible choice for the household if and only if $\zeta_t(a, e) < e(1 - \chi_t(a, e))$, $V_t^D(a, e)$ is continuous wherever it is defined.

Proposition 3 is one of the key supporting results. In general, an increase in continuation utility of the V^D problem, which is a conditional expectation of V , may harm welfare. The reason is, as V^D increases, it makes households more prone to file for bankruptcy. This could else equal reduce q and, by Lemma 1, lower welfare. Proposition 3 shows this increase in V^D due to an increase in continuation utility can be offset via an increase in the filing cost ζ .

¹³Even in this case, it could affect the optimal policy if the planner weights had $\alpha_t(a, e) > 0$ for some state where $V_t^R(a, e)$ was less than $V_t^O(a, e)$ or undefined (implying $a < 0$). For the quantitative work, $\alpha_t(a, e)$ is only positive for $a = 0, e = 1, t = 1$, which rules out this case.

Proposition 3. *Fix some (a, e, t) with $t < T$ and suppose that the continuation utility of the $V_t^D(a, e)$ problem increases due to some policy change. $V_t^D(a, e)$ can be held constant by increasing $\zeta_t(a, e)$.*

In the U.S., a bankruptcy filing affects a household's credit report for roughly ten years. Proposition 4 shows the model's optimal duration is zero. Else equal, setting $\lambda = 0$ (zero duration) increases utility of a bankruptcy household (Lemma 2). However, this can be offset, via Proposition 3, through an increase in ζ . Consequently, the planner need not impose a bad credit record.

Proposition 4. *For any policy with $\lambda_t > 0$, there exists another policy with $\lambda_t = 0$ generating the same planner utility.*

This result relies on full information. Because of it, there are no bad borrowers: Borrowers prone to default are known and the risk is perfectly priced. If some borrowers were more prone to default (say because of heterogeneous default costs like in Athreya, Tam, and Young, 2012b) and this likelihood was private information, there would likely be incentive to exclude bad borrowers from credit markets. The result also relies on partial equilibrium. By imposing $\lambda > 0$, bankrupt households are subject to a no-borrowing constraint. The precautionary savings associated with this is not valued by the planner since it does not affect the aggregate capital stock. In general equilibrium, it would be. However, this result does suggest that, to the extent these assumptions are innocuous, this duration should be shortened.

From this point on, I will consider wlog only policies that have $\lambda_t = 0$ for all t . A convenient property of this is that there is no longer a need to carry around a bankruptcy flag (as households are born with $h = 0$ and never transition to $h = 1$). So, I write $V_t(a, e)$ in place of $V_t(a, e, h = 0)$ and similarly for the policy functions, never referring to $V_t(a, e, h = 1)$ or its policy functions. Another convenient property is that the continuation utilities of the repayment and bankruptcy filing problems are now the same.

As most of the remaining proofs use backward induction to characterize the optimal policy, it is useful to group policies according to whether they agree for future states. The following definition formalizes this notion.

Definition 1. *Two policies are t^* -equivalent if they are identical for all $\tilde{a}, \tilde{e}, \tilde{t} \geq t$ except for one state (a, e, t) .*

Lemma 5 establishes a key method for the planner to improve on an existing policy. Suppose the planner can improve utility V of an age t household while lowering d . In this case, two indirect benefits accrue to age $t - 1$ households. Firstly, they have improved continuation utility. Secondly, they have improved debt pricing. These effects else equal increase V , V^R , and V^D . The increase in V^D can cause default of age $t - 1$ households to increase, which would subsequently worsen credit for $t - 2$ households. However, by holding V^D fixed (Proposition 3), V increases and d decreases for age $t - 1$ households. By repeating the process of holding V^D fixed while letting V and V^R increase for all younger households, welfare is improved.

Lemma 5. *Consider an arbitrary policy and state (a, e, t) . If there is a t^* -equivalent policy such that $V_t(a, e)$ is weakly higher and $\max\{d_t(a, e), o_t(a, e)\}$ is weakly lower, then a policy exists that is weakly better than the original and differs only with respect to ζ_τ for $\tau < t$. If the t^* -equivalent policy has $V_t(a, e)$ strictly higher and $\max\{d_t(a, e), o_t(a, e)\}$ weakly lower, then a strictly better policy exists if $\alpha_t(a, e) > 0$.*

All the costs the planner can inflict on bankruptcy households are deadweight loss. The costs χ and ζ lower utility of filing households without providing any benefit to creditors. But if the planner does not want a household to file, he may simply not let them by setting $D = \{0\}$ (in which case χ and ζ have no effect). If the planner does want a household to file and they do, then the impact on credit is not mitigated by inflicting filing costs. Consequently, the deadweight costs χ and ζ should be set to zero and the planner should simply choose who may or may not file for bankruptcy. Proposition 5 establishes this key result.

Proposition 5. *Any policy specifying $\chi_t(a, e) > 0$ or $\zeta_t(a, e) > 0$ is weakly inferior to a t^* -equivalent policy having $\chi_t(a, e) = \zeta_t(a, e) = 0$. It is strictly inferior if $\alpha_t(a, e) > 0$ and $d_t(a, e) = 1$ under the original policy. Moreover, planner policies may be restricted to $\chi_t(a, e) = \zeta_t(a, e) = 0$ for all a, e, t without loss of generality.*

Propositions 4 and 5 provide the first main result of the paper: All the planner must do is determine who may or may not file for bankruptcy. The planner either allows a household to file for bankruptcy, making them as well off as possible, or the planner completely prevents the household from filing.¹⁴

¹⁴One could easily remove D as a policy instrument and obtain a similar result: $D = \{0\}$ just makes bankruptcy infinitely costly for a household, and the same thing could be accomplished by setting ζ large (i.e., larger than the largest possible endowment).

Since setting $\chi_t(a, e) = \zeta_t(a, e) = \lambda_t = 0$ is weakly optimal, the planner's problem reduces to choosing $D_t(a, e) \in \{\{0\}, \{0, 1\}\}$ for all a, e, t . This is a finite optimization problem, and every choice is feasible (Proposition 1). Consequently, a solution to the planner problem exists, a fact established in Proposition 6.

Proposition 6. *An optimal policy having $\chi_t(a, e) = \zeta_t(a, e) = \lambda_t = 0$ for all a, e, t exists.*

The remaining results exploit equilibrium properties when there are no filing costs. In the absence of filing costs, households can always do as well as autarky by filing for bankruptcy. Since the outside option is assumed to be weakly worse than autarky, $V^D \geq V^O$, as established in Lemma 6.

Lemma 6. *Any policy having $\chi_t(a, e) = \zeta_t(a, e) = 0$ for all a, e, t has $V_t^D(a, e) \geq V_t^O(a, e)$ (for all a, e, t).*

Proposition 7 constitutes the second main result of the paper: The planner allows bankruptcy whenever households cannot repay or would find the outside option preferable. To see why, recall creditors receive nothing when households choose the outside option, $o = 1$, or file for bankruptcy $d = 1$. So, all that matters for credit pricing is $\max\{o, d\}$, as seen in (9). If V^O is preferable to V^R or V^R is undefined, then $V = V^O$ if the planner forbids bankruptcy. If the planner allows bankruptcy, then utility is higher since $V = \max\{V^D, V^O\} = V^D$, the second equality following from Lemma 6. Consequently, the household is made better off while weakly decreasing default. By Lemma 5, this is optimal.

Proposition 7. *Without loss of generality, the optimal policy specifies $D_t(a, e) = \{0, 1\}$ whenever $V_t^O(a, e) > V_t^R(a, e)$ or $V_t^R(a, e)$ undefined. If $V_t^D(a, e) > V_t^A(a, e)$ or $V_t^O(a, e) < V_t^A(a, e)$ and $\alpha_t(a, e) > 0$, then the optimal policy must specify $D_t(a, e) = \{0, 1\}$.*

A trivial consequence of Proposition 7 is that the planner should not set $D = \{0\}$ for all states: Because V^R is undefined if debt is large enough, $D = \{0, 1\}$ is optimal for some highly indebted states. The reason that is interesting is a natural borrowing limit economy, defined in Definition 2, requires $D = \{0\}$ for all states. Consequently, a natural borrowing limit economy is suboptimal, as Corollary 1 states. Note that Definition 2 precisely captures the notion of a natural borrowing limit economy in

Aiyagari (1994): (1) households always choose to repay their debt; (2) they can borrow at a risk-free rate up to the net present value of the worst possible endowment stream; and (3) they never borrow more than this amount.¹⁵ As formally stated in Corollary 1, a natural borrowing limit economy is suboptimal.

Definition 2. *A natural borrowing limit economy is one in which $D_t(a, e) = \{0\}$ for all a, e, t and $V_t^O(a, e)$ is sufficiently low, in the sense of Proposition 2, as to make the outside option irrelevant.*

Corollary 1. *A natural borrowing limit economy is weakly inferior to one allowing bankruptcy whenever $V_t^O(a, e) > V_t^R(a, e)$ or $V_t^R(a, e)$ is undefined. If $V_t^D(a, e) > V_t^A(a, e)$ or $V_t^O(a, e) < V_t^A(a, e)$ and $\alpha_t(a, e) > 0$, then it is strictly inferior.*

This result settles an important question of whether the natural borrowing limit is optimal. This question has been investigated quantitatively by numerous authors who find that implementing a natural borrowing limit greatly improves welfare (relative to benchmark economies calibrated to the U.S.). However, this result shows there is always a better policy allowing default for a wide class of utility functions and endowment processes—even endowment processes where the natural borrowing limit is very large.

Corollary 2. *Wlog, the optimal policy results in the outside option never being chosen.*

Another trivial consequence of Proposition 7 is stated in Corollary 2: The optimal policy has the outside option never chosen in equilibrium. While the outside option is never chosen, it still plays a critical role. Specifically, it pins down in which states of the world the planner “must” allow a household to default. The planner allows default whenever $V^R < V^O$. If V^O is the utility from autarky, this occurs for small amounts of debt. If V^O is very negative, this only occurs for large levels of debt. In fact, Proposition 8 shows that if the outside option value is equal to autarky, then no debt can be supported in equilibrium. Consequently, the best the planner can do is also the worst he can do.

Proposition 8. *If $V_t^O(a, e) = V_t^A(a, e)$ for all a, e, t , then an optimum features $q_t(a', e) = 0$ for all $a' < 0$ and all e, t . That is, an optimum features a zero borrowing limit. Moreover, the maximum of the planner’s problem is also the minimum.*

¹⁵While not expressly forbidden from borrowing more than this, they choose not to as Proposition 2 shows.

4 Quantitative Results

The theoretical results show the planner problem may be reduced to choosing who or who may not file for bankruptcy, via $D_t(a, e)$, and giving bankrupts a fresh start, $\chi_t(a, e) = \zeta_t(a, e) = \lambda_t = 0$ for all a, e, t . Because of this, households can be classified into three groups according to their value of repayment. The first has $V_t^R(a, e)$ less than $V_t^O(a, e)$ or undefined. By Proposition 7, the planner specifies $D_t(a, e) = \{0, 1\}$ and the household chooses $d_t(a, e) = 1$. The second has $V_t^R(a, e)$ greater than $V_t^D(a, e)$. These households choose $d_t(a, e) = 0$ since repayment is always feasible and, in this case, it is the best option. The last has $V_t^R(a, e)$ greater than $V_t^O(a, e)$ but less than $V_t^D(a, e)$. These households will default if and only if the planner lets them. That is, they choose $d_t(a, e) = \max D_t(a, e)$.¹⁶

For this last group, the theory is incomplete. The planner can effectively choose who and who does not default, but it is unclear what he should do. Additionally, the theory does not indicate, in general, how households behave in the optimal policy. In order to better characterize the planner's optimal policy and households' optimal responses, the model is now brought to the data.

4.1 Calibration

To simplify computation, the baseline model period is taken to be three years (as in Livshits et al., 2007). The risk-free price $\bar{q}$ is set to give a 2% annual real interest rate. The discount factor β is set to .96³. The utility function is chosen to be $c^{1-\sigma}/(1-\sigma)$ with $\sigma = 2$.

Households age 1 (real age 23) to $R-1$ (real age 61) have one endowment process, representing working age earnings risk, and households aged R (real age 64) to T (real age 85) have another, representing guaranteed retirement income. For working age households, the endowment follows a discretized version of

$$\begin{aligned} e_t &= \nu \exp(\mu_t + s_t) \\ s_t &= \rho s_{t-1} + \sigma \varepsilon_t, \quad s_1 = 0, \quad \varepsilon_t \sim N(0, 1). \end{aligned} \tag{11}$$

where μ_t is a deterministic age-earnings profile and ν is a normalization constant set so the population-wide average of e is one. The values for μ_t , ρ , and σ are taken

¹⁶These are the only cases to consider because, by Lemma 6, $V_t^D(a, e) \geq V_t^O(a, e)$.

from [Storesletten, Telmer, and Yaron \(2004\)](#).¹⁷ Their estimates, which are annual, are converted to the triennial values $\rho = .952^3$ and $\sigma = .168 * (1 + .952^2 + .952^4)^{1/2}$. Following [Livshits et al. \(2007\)](#), the endowment process for retired households is

$$e_t = .35 + .30e_{R-1}. \quad (12)$$

For welfare comparisons, I specify an “ex-ante” welfare function where $\alpha_1(a = 0, e) = \pi_e$ for all e and $\alpha_t(a, e) = 0$ for all other a, e, t . This has been the most common choice in the literature for lifecycle models featuring default.¹⁸ This welfare measure builds in the welfare of older households because $V_1(a = 0, e)$ implicitly values lifetime utility at all future states reached with positive probability.

For comparison with the optimal policy, I attempt to capture the current U.S. bankruptcy system. In the U.S., $D_t(a, e) = \{0, 1\}$ for all households (i.e., any household in good standing can file for bankruptcy), so that is what I use. So that the average duration of a bad credit record is 10 years, λ_t is set to .7 (for all t). Following [Gordon \(2013\)](#), I use $\chi_t(a, e) = \chi^0 - \chi^1/e$, choosing χ^0 and χ^1 to match the debt-endowment ratio and percent of households filing. The target values for these statistics are from [Livshits et al. \(2007\)](#).¹⁹ I set $\zeta_t(a, e) = 0$ since the χ^1 term implicitly captures any filing cost.²⁰ The calibrated parameters and moments are given in Table 1.

All that is left is to set the value of the outside option. As a benchmark, I take V^O low enough in the sense of Proposition 2 as to make it irrelevant. Proposition 8 shows the opposite extreme of $V^O = V^A$ results in no borrowing. A middle ground is considered in Section 4.5.

¹⁷They allow the variance of persistent shocks to vary from expansion to recession, so I take the average of the two estimated standard deviations, $\sigma_\eta = .168$. Also, the average earnings profile estimates have age , age^2 , and age^3 as regressors, but this appears to be an error. The correct values appear to be for $age/10$, $(age/10)^2$, and $(age/10)^3$.

¹⁸For instance, [Livshits et al. \(2007\)](#), [Athreya et al. \(2009a\)](#), and [Gordon \(2013\)](#) all use this. While trivial to test other welfare functions, this has not been done for brevity.

¹⁹The annual filing rate of 0.84% (the average Chapter 7, nonbusiness, household filing rate from 1995 to 1999) is converted to a triennial rate of 2.52%. The annual debt-income ratio is converted from .084 to .028. The debt measure is revolving consumer credit. Because of asset exemptions and because certain debts cannot be discharged (such as student loans), there is not one appropriate measure of debt. The [Livshits et al. \(2007\)](#) measure I adopt probably overstates the amount of debt as a in the model is conceptually close to net worth. [Chatterjee et al. \(2007\)](#) target the net worth measures, which are an order of magnitude smaller.

²⁰ Official Chapter 7 bankruptcy filing fees in 2012 are \$306 dollars. This value could be more if one included attorney fees.

Statistic	Data	Model	Parameter	Value
Percent of households filing	2.52	2.55	χ_0	0.26
Debt-Endowment Ratio $\times 100$	2.8	2.8	χ_1	0.08
Bad credit record duration (years)	10	10	λ	0.7
Annual real interest rate	2%	2%	$\bar{q}$	0.942

Table 1: Calibration

4.2 Baseline cases

Before computing the optimal policy, it is useful to consider some baseline cases. The first, $D_t(a, e) = \{0, 1\}$ for all a, e, t , provides a lower bound on utility because it results in $q_t(a, e) = 0$ for all $a < 0, e, t$.²¹ This economy is referred to as the zero borrowing limit (ZBL) economy. Another baseline is the current U.S. system (referred to as US). In it, $D_t(a, e) = \{0, 1\}$ for all states, but $\lambda_t > 0$ and $\chi_t(a, e) > 0$ for some states (see the calibration section for details). A third baseline case, a natural borrowing limit (NBL) economy has $D_t(a, e) = \{0\}$ for all states. The literature has shown repeatedly that a NBL drastically improves welfare relative to the US economy in the environment considered in this paper.²² The last baseline must, by Corollary 1, improve on the NBL economy. This policy specifies $D_t(a, e) = \{0\}$ for all states having $V_t^R(a, e) < V_t^O(a, e)$ or $V_t^R(a, e)$ undefined. I refer to this as the $D = 0$ economy, but it should not be confused with $D_t(a, e) = 0$ for all a, e, t , which is the NBL economy.

The results for ZBL, NBL, US, and $D = 0$ are summarized in Table 2. The reported welfare gain is the consumption equivalent welfare measure relative to ZBL.²³ The current US system improves on the worst case scenario, ZBL, with a welfare gain of 1.3%, but the NBL does much better with a welfare gain of 8.6%. Comparing the amounts of debt in US and NBL, one sees a roughly ten-fold increase in NBL. It should be noted at this point that the lowest possible endowment realization is artificially large in this economy (it is just three standard deviations below the mean for persistent earnings shocks and there are no transitory shocks). Consequently, the reported welfare gain and debt statistics of the NBL economy are exaggerated. However, this finding agrees with a large literature that has found the NBL does

²¹While I don't prove this, one could prove it using the reasoning in Proposition 8. At any rate, one can just think of this as setting $V^O = V^A$.

²²I.e., an environment which features a minimum endowment level that is significantly positive.

²³In other words, it is the percent increase in consumption required in every state to make a household indifferent between staying in ZBL with the consumption increase and moving to an alternate economy.

much better than the current U.S. system (in this environment).

Statistic	ZBL	US	NBL	$D = 0$	D^*	Simple
Welfare gain	0.00	1.32	8.55	9.88	12.83	12.23
Total filings (%)	0.00	2.55	0.00	8.06	10.89	7.35
Total debt	0.00	0.03	0.33	0.70	0.82	0.69
Charge-off rate (%)	-	12.9	0.0	50.0	30.5	31.5
Pop. in debt (%)	0.0	26.2	52.8	42.4	50.6	44.4
Total assets [†]	0.93	0.87	0.40	0.31	-0.12	0.03
Total consumption	1.05	1.05	1.02	1.02	0.99	1.00
Avg. interest rate (%)	-	19.0	6.1	432.4	85.8	158.8

[†]Measured as $\int a(1 - d)d\mu$ (capital in general equilibrium).

Table 2: Welfare and Allocations from Differing Default Policies

The $D = 0$ case results in a welfare gain that is roughly 1.3 percentage points larger than the NBL economy. Corollary 1 shows the NBL economy must be weakly inferior to $D = 0$, and here it is strictly suboptimal. The two policies are fairly similar in that they both forbid bankruptcy whenever a household has the ability to repay their debt. Where $D = 0$ and NBL differ is when the household cannot repay, in which case $D = 0$ allows the household a fresh start and NBL does not. While implementing NBL would be very costly (probably involving a return to debtors' prisons), implementing $D = 0$ would be much less costly: If a household filed for bankruptcy, one would determine whether they could repay, and, if not, give them a fresh start.

As reflected in the average interest rate, a large amount of what could be called moral hazard is embedded in the $D = 0$ economy. In particular, some households can easily file for bankruptcy *next* period by entering into an agreement with a creditor to repay millions of dollars, receiving a trifle for that pledge, and entering next period with a massive amount of debt and no way to repay. This strategy is only possible for households in some states. Loosely speaking, it is only possible if $e + a \geq 0$ because this strategy requires $c = e + a - q_t(a', e)a' \approx e + a$. Consequently, if $-a < e$ (i.e., the debt-endowment ratio is less than one), the household cannot implement this strategy. Additionally, they are only allowed to file for bankruptcy in the *current* period if $e + a + \max_{a'}(-q_t(a', e)a') \leq 0$ because no feasible consumption exists in this case. Hence, in the range $-a \in (e, e + \max_{a'}(-q_t(a', e)a'))$, they cannot file for

bankruptcy in the current or future period.²⁴

4.3 The optimal bankruptcy rule

The full optimal policy, labeled as D^* , is now considered. It is computed using a genetic algorithm, and interested readers may consult Appendix B for more details. Some summary statistics are presented in Table 2. D^* generates a 12.8% welfare gain relative to ZBL, more than 11 percentage points higher than US and more than 4 higher than NBL. Along with the large welfare gain is a tremendous amount of debt accumulation, reflected in total asset holdings becoming negative, and a default rate roughly four times higher than the current US rate.

Figure 1: Lifecycle Profiles for D^* , NBL, and US

The lifecycle profiles, presented in Figure 1, reveal a bit more about the differences. One can see that despite a hump-shaped average endowment, average consumption

²⁴In the model, creditors are perfectly comfortable with repayment agreements having very high interest rates as they bear no risk (the return of each debt contract is known with certainty). It is unclear whether creditors would be okay with such an arrangement in real life. If not, the bankruptcy law could permit a household to file for bankruptcy “next period” by having a waiting period based on the conditions above.

is actually at its highest for young households in D^* . This is not the case for US and NBL where consumption is at its lowest for young households. Households under D^* finance their early consumption through a tremendous amount of debt accumulation. This is less so under NBL and practically nonexistent under US. Moreover, default is used throughout working life in D^* but primarily by the very young in US.²⁵ A composition effect is at work here: Middle-aged households have a large negative net worth under D^* , leading many to default, but a slight positive net worth under US.

The optimal policy is represented in Figure 2 for select ages.²⁶ A red dot means that the household has $d_t(a, e) = \max D_t(a, e)$ and that the state is visited in equilibrium (recall from the discussion at the start of Section 4 that households can be classified into three categories, one of which defaults if and only if the planner lets them). A blue asterisk means $\max D_t(a, e) = 1$ and so the household defaults. The vertical axis gives the number of unconditional standard deviations a household's endowment is from the mean for their age. The horizontal axis gives debt (relative to the average endowment which is normalized to one).

While the optimal policy is very complicated, there are at least three discernible patterns which work together to provide the optimal tradeoff between intertemporal consumption smoothing and intratemporal smoothing. First, bankruptcy is usually not allowed for the luckiest households, those two or more standard deviations above mean. By preventing lucky households from filing in some state $a', t + 1$ while permitting unlucky ones to file, $q_t(a', e) = \sum_{e'} \pi_{ee'|t}(1 - d_{t+1}(a', e'))$ can be propped up. This improves intertemporal smoothing in state e, t by lowering consumption for $t + 1$ households with high e' and raising it for low e' households. In a sense, this is a win-win: It reduces both intertemporal variance and intratemporal variance of consumption.

Second, bankruptcy is typically allowed for the largest debt amounts but often not for moderate debts. This feature results in $q_t(a', e)$ being decreasing in debt. As can be seen in Figure 3, which plots price schedule averaged over e, t states for several bankruptcy laws, this effect is fairly small.²⁷ For small debts, D^* results

²⁵Livshits et al. (2007) document filing rates reach a peak at around age 42 whereas in this calibration the peak is much earlier at age 26. The discrepancy is partly due to assuming all households begin life with $a = 0$ and $e = 1$.

²⁶The model begins with age 23, but the policy here does not matter since $a = 0$ for all $t = 1$ households. Only working ages are presented as these are the most important for welfare (young households need access to credit for both insurance and lifecycle reasons).

²⁷The average price schedule is $\sum_{e,t} q_t(a', e) \pi_{et}$ where π_{et} is the invariant distribution over states

Figure 2: The Optimal Policy by Age

in a q of approximately .65 (an annual interest rate of 15%), which declines to .4 (an annual rate of 36%) by the time debt reaches one and a half times the average endowment. By keeping interest rates at around 15% for small debts, households are able to accumulate debt over time, borrowing each period and rolling over old debts. As the accumulated debt becomes large, the probability a household can discharge that debt increases significantly.

The last pattern is essentially a lack of a pattern: The policy is non-monotonic in every variable, a , e , and t . What can the planner hope to achieve in this manner? Figure 3 suggests an answer. In contrast to the other policies, D^* always has q less than the risk-free price because some households are always allowed to default. Given that q is just $\bar{q}$ times the probability of repayment, one can compute the probability of defaulting as $1 - q/\bar{q}$. For a debt level of one, that probability of default is around 35%. It rises to around 70% for a debt level of two. Consequently, for any level of debt, the odds that a household will not be able to default at some point soon, are very small. With 51% of the population in debt and a filing rate of 11%, about one in five indebted households file for bankruptcy each period.

Any attempt to implement the optimal rule in reality would be a complete disaster.

(e, t) .

Figure 3: Average Price Schedules for Various Bankruptcy Laws

It requires that households and creditors have intricate knowledge about not just endowment processes, but about every aspect of bankruptcy law. While this may be feasible for creditors, it is not for households. However, it may be the case that certain aspects of the optimal policy are relatively unimportant. For instance, it may be the case that the non-monotonicity of the optimal policy improves only slightly on a far simpler, monotone policy. The next section looks for and finds such a rule.

4.4 A good simple rule

While D^* , by definition, achieves the best balance between intertemporal and intratemporal consumption smoothing, it is a complicated rule. In the model, households and creditors perfectly understand the structure of D^* and respond accordingly. In practice, such a policy would be far too difficult to implement and would likely be viewed as arbitrary and unfair. In this section, I consider simpler policies that sacrifice welfare gains for ease of implementation.

Recall the optimal policy was characterized, in part, by preventing lucky households from filing and by allowing bankruptcy for large levels of debt. These observa-

tions suggests a rule of the following form,

$$D_t(a, e) = \begin{cases} \{0, 1\} & \text{if } -a/e \geq \gamma \\ \{0\} & \text{if } -a/e < \gamma \end{cases} \quad (13)$$

where γ is a positive constant, could do well for a suitably chosen γ .²⁸ Such a rule has a clear interpretation: A household is allowed to file for bankruptcy if and only if their debt-endowment ratio is greater than γ . As it turns out, such a rule does extremely well. Given the parsimonious functional form in (13), solving for the best such simple rule is just a one-dimensional maximization problem (details are provided in Appendix B).

The best simple rule specifies bankruptcy is allowed if the debt-endowment ratio is greater than 2.64 (i.e., $\gamma = 2.64$). Summary statistics are presented in Table 2 under the heading “Simple.” The welfare gain relative to ZBL is 12.2%, about 3.5 percentage points better than NBL and only .6 points worse than D^* . The filing rate is considerably lower than the D^* rate, although interest rates are somewhat higher.

Figure 4: The Simple Rule by Age

The simple rule eliminates the non-monotonic behavior of the optimal policy while

²⁸Rules specifying a cutoff in the endowment, or a cutoff in debt levels, or of a form $1[.5 < \gamma_0 + \gamma_1 a + \gamma_2 e + \gamma_3 t]$ did not fair as well.

capturing its other aspects. This is clearly seen when comparing Figure 4, which plots the $D_t(a, e)$ implied by the simple rule, and Figure 2, which plots the optimal rule. In particular, bankruptcy is only allowed at the most extreme debt levels and forbidden for lucky households. Because the simple rule is monotone, it also results in much more conventional looking price schedules, as seen in Figure 3. The average price schedule strikes a balance between those implied by $D = 0$ and D^* .

4.5 The impact of a larger outside option

Until now, V^O has been set low enough to be irrelevant leaving households with, essentially, only one default option, bankruptcy. The theory shows the opposite extreme where $V^O = V^A$ completely restricts the planner so that the best he can achieve is also the worst, a zero borrowing limit economy. In reality, V^O is not low enough to be irrelevant nor as large as V^A : If households formally default using Chapter 13 bankruptcy, they have to agree to a repayment plan; if households informally default by refusing to pay, they may be subject to wage garnishment.

As a middle ground, I consider a value for the outside option equal to consuming the endowment each period, i.e. $V_t^O(a, e) = u(e) + \beta \sum_{e'} \pi_{ee'|t} V_{t+1}^O(a, e')$ with $V_{T+1}^O := 0$. For this outside option, it turns out that $D = 0$ does as well as D^* . Each of these produces a welfare gain relative to ZBL of 5.3%, substantially less than the 9.9% ($D = 0$) and 12.8% (D^*) gains when V^O was very low. Additionally, the best simple rule now does worse than $D = 0$, resulting in a welfare gain of only 4.7%. The best debt-endowment cutoff falls from 2.64 to 0.62.

A large outside option severely hampers the planner's ability to maximize welfare. Proposition 8 shows in the extreme of $V^O = V^A$, the planner is incapable of doing better than ZBL. When V^O is lower, equal to consuming the endowment each period, the planner can do substantially more, producing a welfare gain of 5%. However, this gain is paltry relative to the almost 13% gain the planner can achieve when households have, essentially, no outside option.

5 Conclusion

The optimal bankruptcy law offers a fresh start, but this offer is not made to all households. While it always permits bankruptcy when a household cannot repay or

would prefer the outside option, bankruptcy is permitted in other states on a case-by-case basis. The law tends to prevent lucky households from filing and to provide bankruptcy for the largest debt amount. It also allows bankruptcy in other states in a non-monotonic, seemingly haphazard fashion. These features all work together to provide maximal consumption smoothing to households, but not all the features are equally important. A simple cutoff rule in the debt-endowment ratio comes close to the optimal rule but is monotone. As the value of the outside option increases, bankruptcy law becomes impotent.

To characterize the optimal policy theoretically, the current analysis made three assumptions that should be relaxed in future research. First, full information was assumed, which implied there were no “bad borrowers” and hence no incentive to exclude households from credit markets. In a partial information setting, like in [Athreya et al. \(2012b\)](#), this would not be the case. Second, debts were assumed to be completely unsecured with bankruptcy costs resulting in zero payments to creditors. Since the costs were deadweight loss, the planner found it optimal to exclude them completely. If the costs born by households could be transferred to creditors, this might not be the case. Third, partial equilibrium was assumed, making interest rates and the endowment invariant to changes in policy. In a general equilibrium model, policies with a tremendous amount of borrowing would depress wages and raise interest rates. Consequently, it is likely the planner would find such policies suboptimal.²⁹ While relaxing any of these assumptions will likely make theoretical results difficult to obtain, the insights from this paper can be used to guide future quantitative investigations.

References

- S. R. Aiyagari. Uninsured idiosyncratic risk and aggregate savings. *Quarterly Journal of Economics*, 109(3):659–684, 1994.
- K. B. Athreya, X. S. Tam, and E. R. Young. Unsecured credit markets are not insurance markets. *Journal of Monetary Economics*, 56(1):83–103, 2009a.

²⁹Relatedly, the model abstracts from entrepreneurship. Given the individual risks and economy-wide benefits associated with entrepreneurship, the planner may want to allow bankruptcy in more states of the world. [Meh and Terajima \(2008\)](#) investigate the relationship between entrepreneurship and bankruptcy law.

- K. B. Athreya, X. S. Tam, and E. R. Young. Are harsh penalties for default really better? Working Paper 09-11, Federal Reserve Bank of Richmond, 2009b.
- K. B. Athreya, J. M. Sanchez, X. S. Tam, and E. R. Young. Bankruptcy and delinquency in a model of unsecured debt. Working Paper 2012-042C, Federal Reserve Bank of Saint Louis, 2012a.
- K. B. Athreya, X. S. Tam, and E. R. Young. A quantitative theory of information and unsecured credit. *American Economic Journal: Macroeconomics*, 4(3):153–183, 2012b.
- K. B. Athreya, J. M. Sanchez, X. S. Tam, and E. R. Young. Labor market upheaval, default regulations, and consumer debt. Working Paper 2014-002A, Federal Reserve Bank of Saint Louis, 2014.
- S. Chatterjee and G. Gordon. Dealing with consumer default: Bankruptcy vs garnishment. *Journal of Monetary Economics*, 59(S):S1–S16, 2012.
- S. Chatterjee, D. Corbae, M. Nakajima, and J.-V. Ríos-Rull. A quantitative theory of unsecured consumer credit with risk of default. *Econometrica*, 75(6):1525–1589, 2007.
- D. Chen. The impact of personal bankruptcy on labor supply decisions. Mimeo, 2013.
- D. Chen and D. Corbae. On welfare implications of restricted bankruptcy information. *Journal of Macroeconomics*, 33(1):4–13, 2011.
- P. J. Coleman. *Debtors and Creditors in America: Insolvency, Imprisonment for Debt, and Bankruptcy, 1607-1900*. The State Historical Society of Wisconsin, Madison, Wisconsin, 1974.
- J. Dávila, J. H. Hong, P. Krusell, and J.-V. Ríos-Rull. Constrained efficiency in the neoclassical growth model with uninsurable idiosyncratic shocks. *Econometrica*, 80(6):2431–2467, 2012.
- G. Gordon. Evaluating default policy: The business cycle matters. Mimeo, 2013.
- B. Grochulski. Optimal personal bankruptcy design under moral hazard. *Review of Economic Dynamics*, 13(2):350–378, 2010.

- W. Li and P. Sarte. U.S. consumer bankruptcy choice: The importance of general equilibrium effects. *Journal of Monetary Economics*, 53(3):613–631, 2006.
- W. Li and M. White. Mortgage default, foreclosure and bankruptcy. Working Paper 15472, NBER, 2009.
- W. Li, M. J. White, and N. Zhu. Did bankruptcy reform cause mortgage defaults to rise. *American Economic Journal: Economic Policy*, 3(4):123–147, 2011.
- I. Livshits, J. MacGee, and M. Tertilt. Consumer bankruptcy: a fresh start. Working Paper 617, Federal Reserve Bank of Minneapolis, 2003.
- I. Livshits, J. MacGee, and M. Tertilt. Consumer bankruptcy: a fresh start. *American Economic Review*, 97(1):402–418, 2007.
- C. A. Meh and Y. Terajima. Unsecured debt, consumer bankruptcy, and small business. Working Paper 08-5, Bank of Canada, 2008.
- K. Mitman. Macroeconomic effects of bankruptcy and foreclosure policies. PIER Working Paper 11-015, University of Pennsylvania, 2011.
- M. A. Robe, E.-M. Steiger, and P.-A. Michel. Penalties and optimality in financial contracts: Taking stock. Working Paper 2006-013, Humboldt-Universität zu Berlin, 2006.
- K. Storesletten, C. Telmer, and A. Yaron. Cyclical dynamics in idiosyncratic labor market risk. *Journal of Political Economy*, 112(3):695–717, 2004.
- H.-J. Wang and M. J. White. An optimal personal bankruptcy procedure and proposed reforms. *Journal of Legal Studies*, 29(1):255–286, 2000.
- M. J. White. Bankruptcy reform and credit cards. *Journal of Economic Perspectives*, 21(4):179–199, 2007.

A Appendix: Proofs

Proof of Proposition 1. Consider $t = T$ and any a, e, h . Continuation utilities are zero and households solve static problems taking $q_T := 0$ as given. If $h = 0$,

then the outside option is a feasible choice. If $h = 1$, then $a' = 0$ is a feasible choice for any $a \geq 0$ (recall $V_t(a, e, h = 1)$ is only defined for $a \geq 0$). Since, there are only a finite number of choices since a' , d , and o and the choice set is non-empty, a maximum exists. So, $V_T(a, e, h)$ is well-defined for all a, e, h . Moreover, the maximization problem delivers decisions $d_T(a, e)$ and $o_T(a, e)$. Since a and e were arbitrary, $q_{T-1}(a', e)$ is well-defined for all a', e .

Now, consider $t < T$. For induction, assume V_{t+1} and q_t are well-defined. Consider an arbitrary a, e, h . If $h = 0$, taking the outside option is a feasible choice. If $h = 1$, choosing $a' = 0$ is a feasible option. Since there are only a finite number of choices of a' , d , and o , a maximum exists. A solution to the maximization problem gives decisions $d_t(a, e)$ and $o_t(a, e)$. Since a and e were arbitrary, $q_{t-1}(a', e)$ is well-defined for all a', e . This completes the induction argument. $\square$

Proof of Lemma 1. A larger $q_t(a', e)$ for $a' \leq 0$ increases the budget constraint of a household choosing to repay their debt. Since u is increasing, this weakly increases $V_t^R(a, e)$ (for all a). $\square$

Proof of Lemma 2. Any feasible policy for a household of type $(a, e, h = 1, t)$ can also be chosen by a household with type $(a, e, t, h = 0)$ by restricting themselves to $a' \geq 0$. $\square$

Proof of Lemma 3. Define $\underline{a} = \min A$. For $h = 0$ or $h = 1$, one feasible policy when $a \geq 0$ is to choose $a' = 0$ in all subsequent periods and never default. Since the endowment is bounded below by $\underline{e}$, this policy results in lifetime utility greater than or equal to $\sum_{\tau=t}^T \beta^{\tau-t} u(\underline{e})$. So, $V_t(a, e, h)$ is bounded below. Additionally, the most consumption possible is bounded above by $-\bar{q}\underline{a} + \bar{e}$. So, $V_t(a, e, h)$ is bounded above by $\sum_{\tau=t}^T \beta^{\tau-t} u(-\bar{q}\underline{a} + \bar{e})$. These bounds are functions of t , but taking the max and min over all t gives uniform bounds on V . $\square$

Proof of Proposition 2. For the outside option to be chosen in equilibrium, there must be a state (a, e, t) reached with probability $p > 0$ from an initial state $a_1 = 0, e_1, h_1 = 0$ such that $o(a, e, t) = 1$. Consequently, as viewed by an $(a_1, e_1, t = 1)$ household, the expected utility derived from this state is $p\beta^{t-1}V_t^O(a, e)$.

Let the uniform bounds on V be denoted $\underline{V}$ and $\bar{V}$. Wlog, take $\bar{V}$ to be positive. Then expected utility from all other states, as viewed by an $(a_1, e_1, t = 1, h_1)$ household, is not more than $\sum_{t=1}^T \beta^{t-1} \bar{V}$ (a very loose upper bound). So, $V_{t=1}(a_1, e_1, h_1) \leq$

$\sum_{t=1}^T \beta^{t-1} \bar{V} + p\beta^{t-1} V_t^O(a, e)$. Because $a_1 = 0$, Lemma 3 has $V_{t=1}(a_1, e_1, h_1) \geq \underline{V}$. Consequently, $\underline{V} \leq V_{t=1}(a_1, e_1, h_1) \leq \sum_{t=1}^T \beta^{t-1} \bar{V} + p\beta^{t-1} V_t^O(a, e)$ implying

$$V_t^O(a, e) \geq \frac{\underline{V} - \sum_{t=1}^T \beta^{t-1} \bar{V}}{p\beta^{t-1}}.$$

Since p must be larger than the least likely history of shocks which itself occurs with positive probability, call it $\underline{p}$, we have

$$V_t^O(a, e) \geq \frac{\underline{V} - \sum_{t=1}^T \beta^{t-1} \bar{V}}{\underline{p}\beta^{t-1}} \geq \frac{\underline{V} - \sum_{t=1}^T \beta^{t-1} \bar{V}}{p\beta^{t-1}}$$

since the numerator is negative.

Consequently, for any $\delta \leq (\underline{V} - \sum_{t=1}^T \beta^{t-1} \bar{V}) / (p\beta^{t-1})$, $V_t^O(a, e) < \delta$ implies the outside option is not chosen in equilibrium. $\square$

Proof of Lemma 4. Fix some a, e, t and suppress dependence on it. Assume $e(1 - \chi) > 0$ (otherwise V^D is not defined and bankruptcy is not a feasible choice for the household).

Consider $t = T$. In this case, we know the optimal a' choice is always zero, and so $V^D = u(e(1 - \chi) - \zeta)$. Since u is continuous, V^D is continuous in ζ (where defined).

Now consider $t < T$. The budget constraint of a household filing for bankruptcy is

$$c + \bar{q}a' + \zeta = e(1 - \chi), \quad c > 0, a' \geq 0.$$

Each ζ corresponds to a set of a' that is feasible. Call this $A(\zeta)$. This set is weakly decreasing in ζ . It is explicitly given by

$$A(\zeta) = A \cap \left[0, \frac{-\zeta + e(1 - \chi)}{\bar{q}} \right)$$

Abusing notation, define

$$V^D(\zeta) = \max_{a' \in A(\zeta)} u(-\bar{q}a' + e(1 - \chi) - \zeta) + W(a')$$

where W is the continuation utility.

To show V^D is left-continuous in ζ , it suffices to show that for a monotonic sequence $\{\zeta_n\}$ with $\zeta_n \uparrow \zeta$ one has $\lim V^D(\zeta_n) = V^D(\zeta)$. In this case, $A(\zeta) \subset A(\zeta_n)$ for

all n . There are two cases to consider.

First, $A(\zeta) = \lim A(\zeta_n)$. Then, since the sets are finite, there exists an N such that $n \geq N$ implies $A(\zeta_n) = A(\zeta)$. Consequently, $n \geq N$ implies

$$V^D(\zeta_n) = \max_{a' \in A(\zeta)} u(-\bar{q}a' + e(1 - \chi) - \zeta_n) + W(a')$$

where the choice set is $A(\zeta)$, which equals $A(\zeta_n)$. Since this is the upper envelope of a finite number of continuous functions, and the choice set does not change, $\lim V^D(\zeta_n) = V^D(\zeta)$. In this case, V^D is left-continuous.

Second, $A(\zeta) \subset \lim A(\zeta_n)$ with $A(\zeta) \neq \lim A(\zeta_n)$. The set difference $A(\zeta_n) \setminus A(\zeta)$ equals $\{a \in A \mid a \geq 0 \text{ and } a \in [\frac{-\zeta + e(1-\chi)}{\bar{q}}, \frac{-\zeta_n + e(1-\chi)}{\bar{q}}]\}$. For the infinite intersection to be non-empty, it must be a singleton, $\tilde{a} := (-\zeta + e(1 - \chi))/\bar{q}$. So, there exists an N such that $n \geq N$ implies $A(\zeta_n) = A(\zeta) \cup \tilde{a}$. Then $n \geq N$ implies

$$V^D(\zeta_n) = \max \left\{ u(-\bar{q}\tilde{a} + e(1 - \chi) - \zeta_n) + W(\tilde{a}), \max_{a' \in A(\zeta)} u(-\bar{q}a' + e(1 - \chi) - \zeta_n) + W(a') \right\}.$$

or, using the expression for $\tilde{a}$,

$$V^D(\zeta_n) = \max \left\{ u(\zeta - \zeta_n) + W(\tilde{a}), \max_{a' \in A(\zeta)} u(-\bar{q}a' + e(1 - \chi) - \zeta_n) + W(a') \right\}.$$

The first term in brackets diverges to $-\infty$: As $\zeta_n \uparrow \zeta$, $u(\zeta - \zeta_n) + W(\tilde{a}) \downarrow -\infty$ since $\lim_{c \downarrow 0} u(c) = -\infty$ and W is bounded. The second term in brackets converges to a finite value (because W is bounded and $-\bar{q}a' + e(1 - \chi) - \zeta_n$ is bounded away from zero).³⁰ Consequently, there exists an $M \geq N$ such that $n \geq M$ implies

$$V^D(\zeta_n) = \max_{a' \in A(\zeta)} u(-\bar{q}a' + e(1 - \chi) - \zeta_n) + W(a').$$

which converges to $V^D(\zeta)$. So, V^D is left-continuous.

Now, to show that V^D is right-continuous in ζ , it is sufficient to show that for a monotonic decreasing sequence $\{\zeta_n\}$ with $\zeta_n \downarrow \zeta$, one has $\lim V^D(\zeta_n) = V^D(\zeta)$. In this case, $A(\zeta_n) \subset A(\zeta)$ for all n . The set difference $A(\zeta) \setminus A(\zeta_n)$ equals $\{a \in A \mid a \geq 0 \text{ and } a \in [\frac{-\zeta_n + e(1-\chi)}{\bar{q}}, \frac{-\zeta + e(1-\chi)}{\bar{q}}]\}$. Because the infinite intersection of $[\frac{-\zeta_n + e(1-\chi)}{\bar{q}}, \frac{-\zeta + e(1-\chi)}{\bar{q}}]$ as $\zeta_n \downarrow \zeta$ is empty, $\lim A(\zeta_n) = A(\zeta)$. Consequently, there exists an N such that $n \geq N$

³⁰To see the latter claim, note that $-\bar{q}a' + e(1 - \chi) - \zeta_n > -\bar{q}a' + e(1 - \chi) - \zeta > 0$ for all $a' \in A(\zeta)$.

implies $A(\zeta_n) = A(\zeta)$. So, for $n \geq N$,

$$V^D(\zeta_n) = \max_{a' \in A(\zeta)} u(-\bar{q}a' + e(1 - \chi)) - \zeta_n + W(a')$$

which converges to $V^D(\zeta)$. Hence, V^D is right-continuous in ζ . □

Proof of Proposition 3. Fix a, e, t and suppress dependence on it. Let the value of V^D before the policy change be given by V_0^D , and similarly for the continuation utility W and ζ . Then

$$\begin{aligned} V_0^D &= \max_{a' \in A} u(c) + W_0(a') \\ \text{s.t. } c + \bar{q}a' + \zeta_0 &= e(1 - \chi), \quad c > 0, a' \geq 0. \end{aligned}$$

Now, let the new, larger continuation utility be $W_1(a')$. Conditional on ζ , the new value $V^D(\zeta)$ is

$$\begin{aligned} V^D(\zeta) &= \max_{a' \in A} u(c) + W_1(a') \\ \text{s.t. } c + \bar{q}a' + \zeta &= e(1 - \chi), \quad c > 0, a' \geq 0. \end{aligned}$$

It is evident then that $V^D(\zeta_0) \geq V_0^D$ since $W_1 \geq W_0$. If there exists a $\bar{\zeta}$ such that $V^D(\bar{\zeta}) \leq V_0^D$, then the Intermediate Value Theorem shows there exists a $\zeta \in [\zeta_0, \bar{\zeta}]$ such that $V^D(\zeta) = V_0^D$ (by the continuity proven in Lemma 4). So, it suffices to show such an $\bar{\zeta}$ exists.

Consider a monotonically increasing sequence of ζ_n such that $\zeta_n < e(1 - \chi)$ and $\zeta_n \uparrow e(1 - \chi)$. For ζ_n sufficiently close to its limit, the only feasible asset choice is $a' = 0$. Consequently, there exists an N such that $n \geq N$ implies

$$V^D(\zeta_n) = u(e(1 - \chi) - \zeta_n) + W_1(0).$$

Since $W_1(0)$ is bounded and $\lim u(e(1 - \chi) - \zeta_n) \downarrow -\infty$, $V^D(\zeta_n) \downarrow -\infty$. Therefore, there exists an $M \geq N$ such that $n \geq M$ implies $V^D(\zeta_n) < V_0^D$. One can take $\bar{\zeta}$ as ζ_M . □

Proof of Proposition 4. Beginning with a policy that has $\lambda_t > 0$ for potentially all t , I will construct a new policy having $\lambda_t = 0$ for an arbitrary t , but resulting in the exact same value and policy functions for households with $h = 0$ (as well as price

schedules).

Consider a policy with $\lambda_t > 0$. Setting $\lambda_t = 0$ has no effect on future value, policy, or price functions. It also has no effect on $V_t^R(a, e)$ for any a, e . By Lemma 2, it does weakly increase $V_t^D(a, e)$ for each a, e . However, Proposition 3 shows this increase can be undone through a commensurate increase in $\zeta_t(a, e)$. Consequently, V_t^R , V_t^D , $V_t(\cdot, \cdot, h = 0)$, and d_t can all be held constant. This also implies all value functions for $\tau < t$ are unchanged, and so the planner is completely indifferent. $\square$

Proof of Lemma 5. If $t = 1$, increasing $V_t(a, e)$ trivially increases the social planner's utility.

Let $W_{\tilde{t}}(\tilde{a}', \tilde{e})$ be the continuation utility conditional on being in state $\tilde{e}, \tilde{t}$ and choosing $\tilde{a}'$. If $t = 2$, then since $\max\{d_t(a, e), o_t(a, e)\}$ is weakly decreased, q_{t-1} is weakly increased (for all a, e). Additionally, W_{t-1} is also potentially increased. Hence, V_{t-1}^R , V_{t-1}^D , and V_{t-1} weakly increase. While this may change d_{t-1} , it can only increase V_{t-1} .

If $t \geq 3$, then, else equal, d_{t-1} is potentially changed. However, by Proposition 3, an increase in ζ_{t-1} keeps V_{t-1}^D unchanged. Consequently, V_{t-1}^R weakly increases and V_{t-1}^D is held fixed. This induces weakly less default, i.e. d_{t-1} is decreased. Hence, q_{t-2} can only increase. By Lemma 1, this increases V_{t-2}^R else equal. Additionally, the increase in V_{t-1}^R increases V_{t-1} , and hence W_{t-2} . This would drive up V_{t-2}^R and V_{t-2}^D else equal.

Putting the increase in q_{t-2} together with the increase in W_{t-2} , there must be a weak increase in V_{t-2}^R and V_{t-2}^D . Consequently the exact same situation of $t - 1$ arises (where else equal V^R and V^D increase meaning the default decision is potentially changed), but for $t - 2$. The procedure of increasing ζ_τ to offset increases in V_τ^D but allowing V_τ^R and V_τ to increase can be applied now to $\tau = t - 2$. In fact, the same logic can be applied to all $\tau \leq t - 2$ giving the result. $\square$

Proof of Proposition 5. Pick an arbitrary a, e, t and consider any policy that has $\chi_t(a, e)$ or $\zeta_t(a, e)$ not equal to zero. I will construct a t^* -equivalent policy that is weakly better than the original policy and strictly better if $\alpha_t(a, e) > 0$ and $d_t(a, e) = 1$ under the original policy. Since the choice of a, e, t is arbitrary, one can construct a policy having $\chi_t = \zeta_t = 0$ for all t that is weakly better than the original by applying the result for $t = T$ and all a, e , and working recursively backwards for $t = T - 1, \dots, 1$.

Fix a, e, t and suppress dependence on it. First, consider the case where $V^D \leq V^R$

in the original policy. In this case, the household wishes to repay their debt, which is always a feasible choice. By setting $D = \{0\}$ and $\chi = \zeta = 0$, households still choose to repay their debt and receive exactly the same utility. In this case, the planner's utility is unaffected. The case of V^D undefined can be handled in the exact same way.

Second, consider the case where $V^D > V^R$ or V^R is undefined. Setting $\chi = \zeta = 0$ weakly increases V while leaving $\max\{d, o\}$ unchanged: If $D = \{0, 1\}$, then the original policy had $\max\{d, o\} = 1$ and so does the new one; If $D = \{0\}$, then the household cannot change d (which is zero before and after the policy change) and o is unchanged since V^R and V^O are unaffected. Hence, by Lemma 5, this weakly improves the planner's utility. If $d = 1$ under the original policy, then $D = \{0, 1\}$; moreover, since V^D and V increase strictly while $\max\{d, o\}$ is constant, this strictly improves the planner's utility if $\alpha > 0$ (by Lemma 5). $\square$

Proof of Proposition 6. Any policy is weakly inferior to one having $\chi_t(a, e) = \zeta_t(a, e) = \lambda_t = 0$ for all a, e, t . So, if a maximum exists restricting the choice space to $\chi_t(a, e) = \zeta_t(a, e) = \lambda_t = 0$ for all a, e, t , then a maximum exists. In this case, the planner problem reduces to choosing $D_t(a, e) \in \{\{0\}, \{0, 1\}\}$ for all a, e, t . Since there is only a finite number of choices (and all of them are feasible), a maximum exists. $\square$

Proof of Lemma 6. Fix some a, e, t and suppress dependence on it. Recall that $V^O \leq V^A$ (the value of autarky) by assumption. In the case of χ and ζ uniformly equal to zero, autarky is a feasible plan for a bankrupt household. Therefore, $V^D \geq V^A$ and $V^O \leq V^A$ giving the desired result. $\square$

Proof of Proposition 7. Wlog, consider a policy that has $\chi_t(a, e) = \zeta_t(a, e) = \lambda_t = 0$ for all a, e, t . Consider any a, e, t that has $D_t(a, e) = \{0\}$ where $V_t^O(a, e) > V_t^R(a, e)$ or $V_t^R(a, e)$ is undefined. I will construct a t^* -equivalent policy weakly better than the original one. Since the choice of a, e, t was arbitrary, one can then construct a new policy having $D_t(a, e) = \{0, 1\}$ whenever $V_t^O(a, e) > V_t^R(a, e)$ or $V_t^R(a, e)$ is undefined for all a, e, t by applying the result for $t = T$ and all a, e , and working recursively backwards for $t = T - 1, \dots, 1$.

Fix an a, e, t such that $V_t^O(a, e) > V_t^R(a, e)$ or $V_t^R(a, e)$ is undefined and $D_t(a, e) = \{0\}$ and suppress dependence on it. In this state, the household finds it optimal to choose the outside option, so $\max\{d, o\} = 1$. Since $\chi = \zeta = \lambda = 0$, $V^D \geq V^O$ (Lemma 6). Therefore, changing D from $\{0\}$ to $\{0, 1\}$ causes no change in $\max\{d, o\}$

and weakly increases V from V^O to V^D (the increase is strict if $V^O < V^D$ which is implied by $V^O < V^A$ or $V^D > V^A$). By Lemma 5, this increases planner utility with suitable increases in $\zeta_{\tilde{t}}(\tilde{a}, \tilde{e})$ for possibly all $\tilde{a}, \tilde{e}, \tilde{t} < t$. The increase is strict if $V^O < V^D$ and $\alpha > 0$. $\square$

Proof of Proposition 8. Recall the definition of the autarky value as $V_t^A(a, e) = X_t(0, e)$ where X_t defined in (3) can equivalently be defined as

$$\begin{aligned} X_t(a, e) &= \max_{a' \in A} u(c) + \beta \sum_{e'} \pi_{ee'|t} X_{t+1}(\max\{0, a'\}, e') \\ \text{s.t. } c + \bar{q}a' &= e + a, \quad c > 0, a' \geq 0 \end{aligned} \tag{14}$$

with $X_T(a, e) := u(e + a)$.

Consider $t = T$ and an arbitrary a, e . In this case, $V_T(a, e) = V_T^R(a, e) = X_T(a, e) = u(e + a)$. For $a < 0$, $V_T^A(a, e) = V_T^D(a, e) = X_T(0, e) = u(e) > V_T^R(a, e)$. By Proposition 7, it is then weakly optimal to have $D_T(a, e) = \{0, 1\}$. Then, $d_T(a, e) = 1$ because $V_T^A(a, e) = V_T^D(a, e)$ in this case and households file for bankruptcy when indifferent between bankruptcy and the outside option. Because $d_T(a, e) = 1$ for all $a < 0, e$, $q_{T-1}(a', e) = 0$ for all $a' < 0, e$ and $V_T(a, e) = X_T(\max\{0, a\}, e)$ for all a, e .

For induction, suppose that for some t one has $q_t(a', e) = 0$ for all $a' < 0, e$ and $V_{t+1}(a, e) = X_{t+1}(\max\{0, a\}, e)$ for all a, e . Fix some a, e . Conditional on repaying, a choice of $a' < 0$ results in the same consumption as a choice of $a' = 0$ because both have $q_t(a', e)a' = 0$. Moreover, the continuation utilities are the same: $\sum_{e'} \pi_{ee'|t} V_{t+1}(a', e') = \sum_{e'} \pi_{ee'|t} X_{t+1}(0, e')$ for all $a' \leq 0$. Therefore, the household is indifferent over all $a' \leq 0$. Consequently, the household problem conditional on repayment can be written as

$$\begin{aligned} V_t^R(a, e) &= \max_{a' \in A} u(c) + \beta \sum_{e'} \pi_{ee'|t} V_{t+1}(a', e') \\ \text{s.t. } c + \bar{q}a' &= e + a, \quad c > 0, a' \geq 0 \end{aligned}$$

where the choice set is restricted to $a' \geq 0$. Since $V_{t+1}(a', e') = X_{t+1}(\max\{0, a'\}, e')$, it is clear from comparison with (14) that $V_t^R(a, e) = X_t(a, e)$ for all $a \geq 0$ which then implies $V_t(a, e) = X_t(a, e)$ for all $a \geq 0$.

Moreover, the problem conditional on default is

$$\begin{aligned} V_t^D(a, e) &= \max_{a' \in A} u(c) + \beta \sum_{e'} \pi_{ee'|t} V_{t+1}(a', e') \\ \text{s.t. } &c + \bar{q}a' = e, \quad c > 0, a' \geq 0 \end{aligned}$$

Again, it is clear from comparison with (14) that $V_t^D(a, e) = X_t(0, e)$ for all $a < 0$ (i.e., wherever V^D is defined) since the continuation utilities are the same.

So far, it has been established that $V_t(a, e) = X_t(a, e)$ for all $a \geq 0$ and $V_t^D(a, e) = X_t(0, e)$ for all $a < 0$. From the definition of the autarky value, $V_t^A(a, e) = X_t(0, e)$ and so $V_t^D(a, e) = V_t^A(a, e)$. By Proposition 7 it is then weakly optimal to have $D_t(a, e) = \{0, 1\}$ in which case $d_t(a, e) = 1$. This implies $V_t(a, e) = V_t^D(a, e) = X_t(0, e)$ for all $a < 0, e$. Therefore $V_t(a, e) = X_t(\max\{0, a\}, e)$ for all a, e . Moreover, $q_{t-1}(a, e) = 0$ for all $a < 0, e$ since d_t is uniformly equal to zero. This completes the induction argument.

In an optimal policy, $V_t(a, e) = X_t(\max\{0, a\}, e)$ for all a, e, t . In any policy, the household always has $V_t(a, e) \geq \max\{V_t^O(a, e), V_t^R(a, e)\}$ for $a < 0$ and $V_t(a, e) = V_t^R(a, e)$ for $a \geq 0$. Since $V_t^O(a, e) = X_t(0, e)$ for $a < 0$ and $V_t^R(a, e) \geq X_t(a, e)$ for $a \geq 0$, it must be that $V_t(a, e) \geq X_t(\max\{0, a\}, e)$. Therefore, any policy—including the worst policy—does as well as the best policy. $\square$

B Appendix: Computation

This appendix outlines some of the more important parts of the computation, including the algorithms used to compute the optimal policies.

B.1 Discretization

The asset grid is 50 strictly negative points and 250 total points. The points are concentrated near zero. The s shock from (11) is discretized using Tauchen's method with 7 states and a coverage of ± 3 (unconditional) standard deviations.

B.2 Algorithm for finding the full optimal policy

The full optimal policy requires the planner choose $D_t(a, e) \in \{\{0\}, \{0, 1\}\}$ for every $a < 0, e, t$. Consequently, the planner's choice space is very large at $2^{\#A^{--} \times \#E \times T}$ and

an exhaustive search of the choice space is impossible. However, genetic algorithms are designed for optimization problems such as this.

In the language of genetic algorithms, the planner’s policy can be represented as a “gene,” a sequence of 0-1 bits with each bit representing the value of $\max D_t(a, e)$ for one a, e, t . The algorithm works as follows. An initial population of genes is chosen. Over time, each gene mutates and reproduces with other genes in the population. The probability that a given gene reproduces is an increasing function of the gene’s fitness. For the planner problem, I define a gene’s fitness by $-\sum_{a,e,t} \alpha_t(a, e)V_t(a, e)$ since the period utility function is negative.

There are many variants of genetic algorithms that differ in the way genes mutate and reproduce. I use a double crossover variation that generates two children from two parents via cutting the parent at two random positions and exchanging the middle section. Then each child is mutated with the probability of a single bit mutation set to .1%. The total population is set to 16 genes. I use elitism, a variant that saves the best gene by dropping the worst gene at each iteration and replacing it with the best one (consequently, the search is hill-climbing). Every 10,000 iterations (16,000 function evaluations), I check whether the best gene’s fitness has improved by more or less than 10^{-7} . If it is less, the algorithm terminates. Otherwise, it continues.

B.3 Algorithm for finding the simple rule

The simple rule is a one-dimensional maximization problem of finding γ to maximize welfare. It is solved by first doing a 1000 point evenly-spaced grid search from zero to ten. Brent’s method is then used to find a local maximum in a region ± 1 from the grid search maximum. Finally a one-dimensional simplex search is used starting from the maximum Brent’s method found with an initial simplex size of one.