

The Persistence of Financial Distress^{*}

Kartik Athreya[†] José Mustre-del-Río[‡] Juan M. Sánchez[§]

November 9, 2017

Abstract

Using recently available proprietary panel data, we show that while many (35%) US consumers experience financial distress at some point in the life cycle, most of the events of financial distress are primarily concentrated in a much smaller proportion of consumers in persistent trouble. Roughly 10% of consumers are distressed for more than a quarter of the life cycle, and less than 10% of borrowers account for half of all distress events. These facts can be largely accounted for in a straightforward extension of a workhorse model of defaultable debt that accommodates a simple form of heterogeneity in time preference but not otherwise.

JEL classification: D60, E21, E44.

Keywords: default, financial distress, consumer credit, credit card debt.

^{*}For comments and suggestions, we thank seminar participants at the 2016 SED Meetings, 2016 NBER Summer Institute, and 2017 Stockman Conference. We thank Hee Sung Kim for outstanding research assistance. All errors are our own. The views expressed herein are those of the authors and should not be attributed to the FRB of Kansas City, Richmond, St. Louis, or the Federal Reserve System.

[†]Federal Reserve Bank of Richmond; e-mail: kartik.athreya@rich.frb.org

[‡]Federal Reserve Bank of Kansas City; e-mail: jose.mustre-del-rio@kc.frb.org

[§]Federal Reserve Bank of St. Louis; e-mail: juan.m.sanchez78@gmail.com.

1 Introduction

What are the empirics of financial distress in the United States, and to what extent can we understand them through the lens of state-of-the-art quantitative models of defaultable debt? The goal of this paper is to answer these two questions. We tackle the first question using newly available proprietary panel data and tackle the second question estimating multiple models of defaultable consumer debt over the life cycle.

Our empirical work provides estimates of both the extensive and intensive margins of financial distress over the life cycle. We show that while many (35%) US consumers experience financial distress at some point in the life cycle, most distress events are primarily accounted by a much smaller proportion of consumers in persistent trouble: distress incidence is nearly double its unconditional rate even a decade after the initial distress event, and roughly 10% of the consumers are distressed for more than a quarter of the life cycle. In turn, just 10% of borrowers account for half of all distress. We then show that these facts can be largely accounted for in a straightforward extension of a workhorse model of defaultable debt that accommodates a simple form of heterogeneity in time preference but not otherwise. Indeed, the workhorse model of unsecured consumer debt without heterogeneity generates very little persistence.

Our focus on the persistence of financial distress arises because it is the latter that provides essential guidance to the appropriate interpretation of the risks of encountering distress over one's lifetime. For example, if financial distress is highly transitory, a given incidence for it over the life cycle would suggest that most or all households face similar risks over their lives. If, on the other hand, financial distress is highly persistent, the same incidence would be disproportionately accounted for by a much smaller number of borrowers who repeatedly, or in a sustained fashion, encounter distress.

Our work contributes in two ways. First, to our knowledge, our work is novel in providing a complete description of the incidence, concentration, and dynamics of financial distress. Our findings are striking: As for incidence, at least a third of all households will experience financial distress over their lifetime. As for dynamics, financial distress is very persistent. For instance, even a decade after experiencing a period of severe delinquency, borrowers are at roughly twice the risk of being severely delinquent as the unconditional person their age. Moreover, we find that this measure of persistence is essentially invariant over the life cycle, despite unconditional persistence having a clear life cycle hump and decline beyond middle age.¹ As for concentration, financial distress is so unequally distributed that the top 10% of borrowers account for fully half of all financial distress events.

The second contribution of this paper is to show that the facts on financial distress are suggestive of persistent differences in household preferences related to the use of credit and debt repayment. Specifically, we demonstrate that one is led fairly naturally to a model that retains most of the features of a baseline model of defaultable consumer debt but allows for a parsimonious form of heterogeneity in terms of household time preference. This makes clear that the risk of financial distress is one that is resolved early in life, with most borrowers knowing that they will face very little risk in life and a much smaller few knowing that they

¹In addition to these facts, and as with the facts on incidence described above, the persistence of financial distress is very similar across all 50 states. Results are discussed in the Appendix.

will face a bleak outlook.

The difficulty that households appear to have in clearing their debt obligations is suggestive of at least a subset facing significant constraints on access to relatively cheap credit. After all, why would any consumer repeatedly be seen in the data with debt in arrears, if not for their inability to refinance those debts or shift them to creditors with whom they are not delinquent?

One motivation for our work, particularly our empirical efforts, is that in recent work, the extreme events of bankruptcy or outright repudiation play an important role in helping discipline the quantitative strength of limited commitment on allocations. Specifically, it is the observable event of personal bankruptcy that provides a main target for the calibration or parameterization of the models. Recent work uses such models to analyze the implications of regulations (especially bankruptcy law) on outcomes. For example, recent reforms like the Bankruptcy Abuse Prevention and Consumer Protection Act of 2005 (BAPCPA), or the effects of competing social insurance policies on credit use have been studied through the lens of what is now a “standard default model” (e.g., Livshits et al. (2007), Chatterjee et al. (2007)). A consistent finding in this work is that debt relief makes credit expensive and so sensitive to borrower circumstances that the overall ability to smooth consumption (and hence ex-ante welfare) is substantially worsened. A caveat is that sudden large shocks that force households to consume, or spend (e.g., legal judgments or uninsured medical expenses), restore the ability of default to provide net benefits in an ex-ante sense. But as noted above, absent clear evidence that the baseline models used in these analyses capture well the time path of overall financial conditions, such as when measured by the presence of financial distress, there is reason for concern about the sensitivity of that finding.

While our analysis suggests the presence of heterogeneity in discounting, we stress that such variation is still a stand-in for a variety of other forces—notably unobserved demands for consumption within the household arising from a variety of sources. The appropriate interpretation of our findings is therefore not that individuals are necessarily widely varying in their personal levels of patience, but rather that a sizable subset of consumers are persistently rendered effectively impatient, potentially by a host of additional factors not modeled here. Future work that allows for more detail on household-level shocks, intra-household bargaining, and other (persistent) within-household resource variation is therefore essential before reaching conclusions that individuals are to be “implicated” in their fates.² Indeed, it is for this reason that we avoid any normative analysis in this paper.

1.1 Related work

Financial distress or household financial “fragility” has received significant attention in recent work and has been the topic of interest with the general public.³ Interest in the ability of the household to shield itself from susceptibility to shocks through the use of financial markets

²Indeed, the important work of Becker and Mulligan (1997) shows the list of deep forces shaping time-preference is long. More specifically, their analysis shows how income, wealth, mortality, addictions, uncertainty, and other variables affect the degree of time preference. Our work underscores the need for empirical work more capable of allowing researchers to unpack the particular circumstances facing households, especially those of the subset whose consumption needs are, evidently, very persistently urgent.

³<http://www.cbsnews.com/news/the-financial-fragility-of-the-american-household/>

is, of course, longstanding. However, recent work has been aided by the arrival of more detailed data on household balance sheets (Lusardi et al. (2011), Lusardi (2011), Jappelli et al. (2013), Ampudia et al. (2016), Brunetti et al. (2016)) and aims to gauge borrowing capacity and resilience to sudden, unforeseen expenditures. Specifically, this work primarily focuses on measuring the ability of households to remain current on incurred debts, as well as the question of how much borrowing the household could feasibly engage in, in a short term period, e.g., 30 days—especially to cover an unforeseen “expense” (as opposed to a change in income, say). A rough summary of this work might be this: A substantial proportion of households in the US as well as in the EU are, by various measures, “fragile” or in—or near—financial distress.

Our work is also clearly related to the far larger body of work concerned with the measurement of liquidity constraints across consumers. Substantively, this work tries to measure the proportion of US households who are liquidity constrained, and therefore, not well-positioned to deal with adverse shocks. These include papers of Jappelli and Pagano (1999), Hall and Mishkin (1982), Zeldes (1989) and others. More recently, Gross and Souleles (2002) use exogenous variation in credit line extensions to gauge the fraction who increase their debt in response (and hence can be viewed as having been constrained). They find (unsurprisingly) that those close to their limits increased borrowing by most but so did even those further away from their credit limit. The consensus seems to be that roughly 20% are “constrained” either in terms of excess sensitivity to income or in terms of how they respond to survey questions. Compared to this previous literature, our study uncovers the persistence of financial distress. This has important implications for welfare analysis and policy design, as we will show.

Our work contributes to the research programs above in two ways. First, to our knowledge, we are the first to focus on the empirical dynamics of consumer financial distress, which we broadly define to be situations in which the household remains susceptible to any deviation of income from its ex-ante expectation. In this sense, our measures are informed by the line of work emphasizing household insurance, particularly Kaplan and Violante (2010), and the “insurance coefficient” approach of Blundell et al. (2008). Our emphasis, relative to the preceding work, is on direct measures of financial conditions that have empirical counterparts.

Second, our work extracts a previously unknown implication from the “standard default model.” We have already noted above that benchmark models of unsecured consumer debt and default over the life cycle do not imply anything near the observed level of persistence. These include models based primarily on those of Livshits et al. (2007) and Athreya (2008). For example, when distress is measured by severe delinquency (i.e., having a debt 120 days or more past due), the model-implied gaps between the unconditional and conditional probabilities of distress over the life cycle are (i) far too small, at only 15 percentage points at the one-year mark, compared to far larger gap in the data of 60 percentage points, and (ii) far too transitory: At even the three-year mark, the model fails completely to generate separation between the conditional and unconditional probabilities of being in financial distress.

Our empirical results, and the inability of workhorse models to account for them, suggest that underlying persistent heterogeneity may be an important force in consumer behavior. We pursue this intuition and demonstrate that an underlying environment in which households may differ systematically from each other in their “type”, as defined by their patience,

allows for much greater explanatory power.

Our emphasis on understanding household financial distress, and our use of the facts of distress to learn about underlying time preference, are both novel. However, our findings inform a larger body of recently emerging work that uses consumer credit to conclude that permanent heterogeneity in time-discounting is an important feature of the data.⁴ Closest of all is the work of Fulford and Schuh (2017), who demonstrate that household credit utilization and life cycle consumption and savings (credit-use) patterns clearly suggest important heterogeneity in time preference. Indeed, these authors estimate that nearly two-thirds (64%) of all households are effectively impatient, enough so to live essentially hand-to-mouth. Our work strongly complements theirs by showing that the facts of financial distress—a state that is clearly and unambiguously observable—drive one to reach very similar conclusions. In particular, our model differs from theirs, and all other previous work, by deriving financial distress from a model that incorporates default as an option for borrowers. This in turn allows our work to capture the complications posed by default risk for credit pricing and availability. Notably, terms across borrowers will vary (both over time for a given borrower and across different borrowers at any given time) in response to the evolution of their balance sheet and future earnings prospects.

Two other recent papers also use credit market data to conclude that there is nontrivial variation in patience across borrowers. First, Gorbachev and Luengo-Prado (2016) use National Longitudinal Survey of the Youth (NLSY) data to conclude—from the observation of variation in individuals in the extent to which they borrow and save simultaneously—that US households vary substantially in time preference. Second, Meier and Sprenger (2017) conclude in favor of discount-rate heterogeneity from data obtained in a field experiment on credit use. Lastly, while not about consumer credit use, Parker (2017) finds that US households are better described as varying systematically in their preferences than in terms of the shocks they receive based on household consumption responses to random variation in receipt of lump-sum cash transfers (arising from stimulus payments during the Great Recession).⁵ Indeed, he argues that the observed lack of consumption smoothing in those data are “associated with a measure of impatience” among other persistent differences. Overall, the fact that credit use data, financial distress data, and data on consumer response to transfer payments, all point to variation in discounting is noteworthy and suggests that this may be a genuine, and genuinely important, form of heterogeneity.

The remainder of the paper is organized as follows. Section 2 defines our preferred measures of financial distress, and provides an empirical analysis of the behavior of these measures. Section 3 then lays out a variant of a standard life cycle model of consumption and defaultable debt that largely accounts for the empirics of financial distress. Section 4 provides the main comparisons of models with data. Section 5 illustrates that settings that

⁴A much larger literature has used data on consumption and income, and sometimes wealth as well, to estimate models that imply preference heterogeneity more generally. These include early work of Lawrance (1991), Gourinchas and Parker (2002) and Cagetti (2003). Other work on the presence of discount-factor heterogeneity includes: Hausman (1979), Samwick (1998), Warner and Pleeter (2001), and Belzil and Hansen (1999). Lastly, see Frederick et al. (2002) for a survey.

⁵Relatedly, Mustre-del Río (2015) finds that persistent employment differences across males in the US cannot be explained by differences in wealth or wages and hence are indicative of persistent differences in the disutility of work.

do not allow for discount-factor heterogeneity simply cannot capture what would appear to be critical features of observed financial distress. Section 6 concludes.

2 Financial Distress in the US

The first goal of this paper is to establish the empirics of financial distress. As indicated above, we exploit recently available account-level panel data from Federal Reserve Bank of New York Consumer Credit Panel/Equifax. These data cover an 18 year window for a large number of account holders.

We focus on individuals with complete credit histories between 1999Q1 to 2017Q2. Additionally, we restrict our attention to the cohort that enters 1999Q1 between the ages of 25-55. Thus, by the end of our observation period the oldest individuals in our sample are 73, while the youngest are 43. Because our model will focus on default and delinquency behavior prior to retirement we further restrict our measurements to individuals through the age of 65.⁶ While our analysis focuses on a specific cohort over a particular time period, we note that our observations on the incidence of financial distress are robust to looking at repeat cross-sections over the 1999Q1-2017Q2 period. Additionally, our observations on the persistence of financial distress do not seem to be driven simply by behavior during and after the Great Recession. We will define an individual to be in financial distress (FD) in a given year if, in that period, they are recorded as having at least one severely delinquent account: an account for which payment is at least 120 days past due. Additional details about our data appear in the Appendix.

To start, consider first the “extensive” margin of distress: How broadly shared an experience is financial distress? Figure 1 takes a life cycle perspective. The left panel shows the fraction of individuals in delinquency. What emerges is central to what follows: financial distress, while relevant for consumers of all ages, is not widespread. The line in the figure begins near 10-14% among the young, and falls below 10% later in life.


As for the “intensive” margin of financial distress, consider the right panel of Figure 1. We see that the fraction of debt in delinquency follows a very similar pattern, with the youngest having the largest fraction of debt in delinquency (e.g., roughly 13% at age 25) with this proportion falling substantially to around 6% by age 55.

Next, to begin assessing the persistence of financial distress, we compare the unconditional probability of falling into delinquency with the conditional probability. Specifically, we condition on the time elapsed since a transit into financial distress by an individual. In Figure 2 we see very clearly just how persistent the state of distress is for US consumers. Conditional on being in distress today, the likelihood of being distressed in six years (the orange dotted line) is nearly triple that of the unconditional rate (the black line) over the entire life cycle. As we show further below, this particular feature will elude the standard model of defaultable consumer debt and will fairly decisively suggest the importance of heterogeneity in individual time preference.

One might also ask whether an alternative “extensive” margin measure might indicate something different. In particular, instead of defining distress to be a situation in which an individual has severely delinquent debt, one could measure the proportion of consumers


⁶In all figures we plot data through age 55 because we measure default up to 10 years in the future.

Figure 1: The Incidence of Financial Distress Over the Life Cycle



Source: See Appendix

Figure 2: The Persistence of Financial Distress Over the Life Cycle (debt)




Source: See Appendix

who have depleted their available credit (e.g, those who have “maxed out” their credit cards). Figure 3 (which excludes those with 0 credit limit) shows that the answers remain very similar: limited borrowing capacity remains a prevalent issue for a small, but far from negligible, group of borrowers, throughout the life cycle, and just as when measured by delinquency, indicates substantial persistence.⁷

We now provide additional detail on the persistence of financial distress (defined unless otherwise indicated by the presence of severely delinquent accounts on a consumer’s balance sheet). One particular point to keep in mind is that the more transitory distress is, the less

⁷In our dataset, on average, about 50% of individuals in delinquency have also depleted their credit.

Figure 3: The Persistence of Financial Distress Over the Life Cycle (credit limit)




Source: See Appendix

one might view it as relevant to household well-being. In particular, one might conclude that highly fleeting distress indicates optimal use by borrowers of the “real option” to force their creditors to implicitly refinance their loans (subject to the costs associated with being severely late on payments).

Figure 4 provides further evidence that distress is not highly fleeting, but instead persistent over long periods. It displays the distribution of time spent (as a proportion of the 18 years in sample) in financial distress for those who have been delinquent at any time during the sample window. It is clear that while it is indeed fleeting for some (roughly 30%), for 70% of consumers, distress is much more routine state of affairs. Indeed, more than 30% of all those who experience financial distress spend at least a quarter of their lives in it!


A second way to evaluate the persistence of financial distress is to examine the number of distinct spells of delinquency that an individual will experience. Figure 5 shows how, for those who have experienced financial distress at least once, the number of spells they experience is often substantial, with roughly a tenth of the sample experiencing four or more spells.

Figure 4: The Duration of Financial Distress (share of life)



Source: See Appendix

Figure 5: The Duration of Financial Distress (spells)




Source: See Appendix

A more direct way of showing that a small share of the population accounts for most of the financial distress is to look at the concentration of financial distress. In Figures 6 and 7, we show the Lorenz curves for our two measures of financial distress. They show that most of the 80% of the financial distress is accounted by less than 20% of people.


The measures presented thus far are “extensive margin” measures: They are based on measures of financial distress that are binary—whether or not someone has severely delinquent debt, or whether or not someone has reached their credit limit. While the data suggest

Figure 6: The Concentration of Financial Distress (debt)



Source: See Appendix

Figure 7: The Concentration of Financial Distress (credit limit)




Source: See Appendix

that by these metrics, financial distress is not only frequent, but also persistent, it might still not be an economically important phenomenon if the debts that consumers are severely delinquent on are themselves trivial. We now demonstrate that they are not.

A natural intensive-margin measure is one that relates the volume of debt in delinquency to the total debt owed by individuals over the life cycle. Figure 8 clearly indicates that financial distress among those facing it is “intense,” measured either in terms of the average (across individuals) proportion of debt or average number of accounts severely delinquent.

When debt is the measure (left panel), we see not only do distressed borrowers have almost all (roughly 80%) their debts in delinquency, but also that there is virtually no life cycle component to the intensity of distress, as the intensity of distress falls from nearly 88% early in the life cycle to 83% at older ages.

Figure 8: The Intensity of Financial Distress Over the Life Cycle



Source: See Appendix

Figure 9 summarizes the amount of delinquent debt by 50th, 75th and 90th percentile by age.⁸ It shows that the median debt in delinquency is fairly stable over the life cycle, but the upper tail (measured either by the 75th or 90th percentiles) grows substantially over the life cycle. Of course, the incidence of distress is lower late in the life cycle, but it is clear that among the distressed, the highest distress occurs among older individuals.

Collecting the extensive- and intensive-margin empirics, we can summarize our findings as follows: Financial distress among US individuals is driven by a relatively small proportion of individuals who experience significant and persistent debt repayment problems. These facts are novel and certainly suggest that financial distress may well be an important phenomenon, especially from the point of view of a subset of individuals looking out over a life cycle.⁹ Discerning this importance, however, requires a model, which we turn to next.


3 Understanding Financial Distress

We now address the question of how well the facts documented above, particularly the persistence of financial distress, can be accounted for in a setting where households make choices over consumption, credit, and repayment. Our model builds on important work of Livshits et al. (2007) because it provides a benchmark life cycle consumption savings model in which debt may be repudiated. Our model will feature two tractable extensions of this

⁸Each year has been deflated by January's Current Price Index for Urban Consumers (CPI-U) as published by the Bureau of Labor Statistics (BLS).

⁹In the Appendix we also establish that the incidence of distress (as measured by severe delinquency) is prevalent in very similar ways across all 50 states. This occurs despite what might seem at first glance to be potentially salient differences in consumer default regulations.

Figure 9: Distressed Debt by Age



Source: See Appendix

environment: (1) Allowance for informal default, and (2) heterogeneity in time preference. We will show that these extensions largely, though not perfectly, account for the facts. We then turn to demonstrating these extensions are, in a sense, “identified” by the data. Specifically, we show that a fairly comprehensive set of alternative models, including ones that allow for income risk to be extremely persistent, and ones that do not distinguish between informal and formal default, each fail to deliver simultaneously the incidence of financial distress, its persistence, or concentration.

3.1 A Benchmark Model

3.1.1 Model

There is a continuum of finitely-lived individuals who are risk-averse and discount the future exponentially. Individuals survive to the next period with probability ϱ_n , which depends on age n . Each agent works for a finite number of periods and then retires at age W . In each period, agents choose consumption c and assets (or debt) a' . Debt may be repudiated in one of two ways. First, the agent may simply cease payment. This is known as delinquency. With delinquency, a household’s debt is not necessarily forgiven, however. Instead, debts are forgiven with probability γ . The probabilistic elimination of debts is meant to capture the presence of creditors periodically giving up on collections efforts. With probability $1 - \gamma$, then, a household’s rolled over debt is not discharged, and in this case, the household pays a “penalty” rate, η , of interest higher than the average rate paid by borrowers.¹⁰ Moreover, in any period of delinquency, consumption equals income up to a threshold τ .¹¹ Second, and as

¹⁰This representation captures the main ingredients in Athreya et al. (2017). They show evidence that penalty rates modeled like this are able to capture key features of delinquency.

¹¹The remaining income is lost, for instance, when dealing with debt collectors.

is standard in models of unsecured debt, agents may invoke formal default via a procedure that represents consumer bankruptcy. If this is the path chosen, all debts are erased, and in the period of filing for bankruptcy, consumption equals income net of the monetary cost f of filing for bankruptcy. Unlike delinquency, there is no income garnishment in bankruptcy.

While all agents are assumed to have identical attitudes toward risk, they will be allowed to vary in their willingness to substitute consumption across time. Specifically, we assume individuals can be divided into two types via their subjective discount factors, with the calibration (described below) determining (1) the magnitudes and (2) the proportion of the population, carrying each of the two values. More precisely, let p_L denote the proportion of individuals who have a discount factor β_L . The remaining $1 - p_L$ share of individuals are potentially more patient and thus have a discount factor $\beta_H \geq \beta_L$. Denote an individual's discount type by j .

In this framework lifetime utility is written as

$$G_{i,j,n}(z, \varepsilon, a) = \max\{V_{i,j,n}(z, \varepsilon, a), B_{i,j,n}(z, \varepsilon), D_{i,j,n}(z, \varepsilon, a)\} \quad (1)$$

where V , B and D are lifetime utilities for households paying back their debt, filing for bankruptcy, and being delinquent on their debt, respectively.

Next, the lifetime utility of bankruptcy is

$$B_{i,j,n}(z, \varepsilon) = u(y_{i,n}(z, \varepsilon) - f) + \varrho_n \beta_j \mathbb{E}[G_{i,j,n+1}(z', \varepsilon', 0)|z]. \quad (2)$$

Recall from above that if bankruptcy is chosen, then in that period, household consumption equals income net of bankruptcy filing costs f , while in the period following bankruptcy, the household has no debt.

Now suppose the household decides to be delinquent on its debt. In this case, lifetime utility reads as:

$$D_{i,j,n}(z, \varepsilon, a) = u(\min\{y_{i,n}(z, \varepsilon), \tau\}) + \varrho_n \beta_j \mathbb{E}[(1 - \gamma)G_{i,j,n+1}(z', \varepsilon', (1 + \eta)a) + \gamma G_{i,j,n+1}(z', \varepsilon', 0)|z]. \quad (3)$$

This reflects the features described above. In particular, it makes clear that in the period of delinquency, household consumption equals income up to a threshold τ , and in the period after choosing to be delinquent, two states can occur: with probability $(1 - \gamma)$ the household's debt is rolled over at an interest rate of η and hence $a' = (1 + \eta)a$. Alternatively, with probability γ the household's debt is fully discharged and hence the household enters the period with no debt (i.e. $a' = 0$).

Finally, suppose the household decides to pay back its debt. This is simply the case of a pure consumption and savings model, with only the continuation value imparting any difference between it and something entirely standard. The consumer who repays debt as promised receives lifetime utility of

$$\begin{aligned} V_{i,j,n}(z, \varepsilon, a) &= \max_{\{a', c\}} u(c) + \varrho_n \beta_j \mathbb{E}[G_{i,j,n+1}(z', \varepsilon', a')|z], \\ \text{subject to} \quad & c + a' q_{i,j,n}(z, a') = a + y_{i,n}(z, \varepsilon), \\ & c \geq 0. \end{aligned} \quad (4)$$

where $q_{i,j,n}(z, a')$ is the price of debt a' and is defined below.

In what follows, the policy function R indicates whether the household pays back its debt (repay), becomes delinquent, or files for bankruptcy:

$$R_{i,j,n}(z, \varepsilon, a) = \begin{cases} 1 & \text{if } V_{i,j,n}(z, \varepsilon, a) = \max\{V_{i,j,n}, B_{i,j,n}, D_{i,j,n}\} \\ 2 & \text{if } D_{i,j,n}(z, \varepsilon, a) = \max\{V_{i,j,n}, B_{i,j,n}, D_{i,j,n}\} \\ 0 & \text{otherwise.} \end{cases}$$

Because default is an option borrowers hold, lenders must be compensated for the risk they bear, at least on average. Specifically, we require that lenders break even in expectation on each loan, given the information they have on borrowers. Specifically, lenders must forecast, based on the borrower's current state, the probability that their income one period hence (when debt comes due) will fall into a set where default (either via delinquency or bankruptcy) becomes more valuable than repayment. The probability of default will, of course depend on the probability distribution of income one period hence, and also on the discount factor of the borrower in question. Let the price of a debt issuance by a given borrower of type i, j, n be given as $q_{i,j,n}(z, a')$. This price function is then taken as given by all borrowers and by virtue of the diversification assumed in the continuum, breaks lenders to such a household type even with probability one. It satisfies the following condition:

$$q_{i,j,n}(z, a') = \frac{1}{1+r} \varrho_n \mathbb{E} \left[\mathbb{I}_{R_{i,j,n+1}(z', \varepsilon', a')=1} + \mathbb{I}_{R_{i,j,n+1}(z', \varepsilon', a')=2} (1-\gamma)(1+\eta) q_{i,j,n+1}(z', a'') | z \right] \quad (5)$$

$$\text{with } a'' = (1+\eta)a'$$

The first term on the right-hand side represents the probability that the household repays its debt. The second term represents the probability that the household chooses to become delinquent when given the option to repay, file for bankruptcy, or become delinquent. This term takes into account that delinquent debt tomorrow is fully discharged at a rate γ .

An equilibrium in this economy is a set of value functions, optimal decision rules for the consumer, default probabilities, and bond prices, such that equations (1) to (4) are satisfied and prices satisfy the zero-profit condition (5).

3.1.2 Calibration

We begin, in Table 1, with the parameters we set externally. The penalty rate for delinquent debt is set to 20% annually, following Livshits et al. (2007). Bankruptcy filing costs are to 2.8% of average income, or roughly \$1,000, again following Livshits et al. (2007). Turning to the income process parameters we follow Kaplan and Violante (2010). During working ages log income equals $y_{i,n} = l_n + z_{i,n} + e_{i,n}$. Here l_n is a life cycle component, $e_{i,n}$ is an i.i.d transitory component, and $z_{i,n}$ is a persistent component that follows a random walk: $z_{i,n} = z_{i,n-1} + \varepsilon_{i,n}$. We assume $\varepsilon_{i,n}$ and $e_{i,n}$ are normally distributed with variances σ_ε^2 and σ_e^2 , respectively. To set retirement income, we follow Hatchondo et al. (2015). While in retirement, the household receives a fraction of the last realization of the persistent component of its working-age income using the replacement ratio formula: $\max\{A_0 + A_1 \exp(z_{i,W-1}), A_2\}$. In order to be consistent with US replacement ratios, we calibrate A_0, A_1 and A_2 such that

the replacement ratio declines with income, from 69 to 14%, with an average replacement rate of 47%. The age-specific survival probabilities follow Kaplan and Violante (2010). Lastly, the initial distribution of wealth and earnings matches wealth holdings for 25- to 26-year-olds from the Panel Study of Income Dynamics (PSID).

Table 1: Benchmark Model Parameters Determined Externally

Parameter	Value
σ , Coefficient of relative risk aversion	2.000
r , Risk free interest rate	3.000%
W , Retirement age	65
η , Roll-over interest rate on delinquent debt	20%
f , Bankruptcy filing cost (as a share of average income)	0.028
σ_ε^2 , Variance of permanent shocks	0.05
σ_e^2 , Variance of transitory shocks	0.01
A_0 , Replacement ratio	0.71
A_1 , Replacement ratio	-0.045
A_2 , Replacement ratio	0.14

In addition to the parameters set without reference to model outcomes, our model features other parameters to be determined in calibration. These are the following: the threshold for income garnishments in delinquency, τ , the probability that delinquent debt is fully discharged, γ , and the parameters governing the distribution of discount factor types, β_L , β_h , and p_L . For simplicity, we anchor the discount factor of the patient borrower at $\beta_H = 1.00$.¹² The remaining parameters are chosen such that the model matches the moments in bold as seen in Table 2. In particular, the model is calibrated to replicate the incidence of financial distress at ages 30 and 50, the average incidence of bankruptcy between the ages of 25 and 65, and the share of individuals who are never in financial distress (over the sample period we observe them).

The implied parameter values are presented in Table 3. In terms of the calibration performance, we see from Table 2 that overall, the baseline model captures well what it targets.

3.1.3 Results from the Benchmark Model

The goal of this part of the paper is to gauge the ability of standard theory to account for the facts of financial distress: Its incidence, its concentration, and its persistence. In Table 2 we see that both in terms of concentration, and in terms of persistence, the model does quite well. The top 20% of model borrowers account for the vast majority (85%) of all financial distress, something very close to what is observed. Similarly, the model captures well, though with a less steady decline and initial level, the persistence of financial distress.

Turning next to a more general summary of concentration, Figure 10 demonstrates excellent congruence between model and data. As we will see below, alternatives to this baseline

¹²Allowing this parameter to be flexibly calibrated does not affect our results as the calibration converges to value very close to 1.0.

Table 2: Benchmark Model Calibration

Moment	Data	Model
FD rate at age 30 (in %)	15.16	15.19
FD rate at age 50 (in %)	10.06	7.54
Average BK rate ages 25-65 (in %)	0.74	0.74
Pop. that never defaults (in %)	64.23	66.09
FD share of top 20% of pop. (in %)	88.07	84.86
Avg. cond. FD rate at 2 years	0.52	0.40
Avg. cond. FD rate at 4 years	0.34	0.33
Avg. cond. FD rate at 6 years	0.23	0.30
Avg. cond. FD rate at 8 years	0.17	0.28
Avg. cond. FD rate at 10 years	0.15	0.26

Note: Numbers in bold denote targeted moments in the calibration of the model.

Table 3: Calibrated Parameters for Benchmark Model

Parameter	Value
Earnings threshold in DQ τ	4.79
Discharge shock to DQ debt γ	0.34
Low discount factor β_L	0.74
High Discount factor β_H	1.00 [†]
Share of pop. of type L	0.56

Note: [†] denotes a parameter fixed by assumption.

model have substantially more difficulty in matching this, and other, aspects of financial distress.

The overall concentration of financial distress across households is one dimension of interest, but it does not directly inform us on the persistence of financial distress over the life cycle. To understand the extent to which the baseline setting predicts how long distress lasts, in probabilistic terms, consider the following findings. In Figure 11 we see that the model captures well the near-term persistence of financial distress, measured here by the conditional probability of being in distress at various leads beyond distress at an initial age, across the entire life cycle. It also largely, but not perfectly captures persistence at most ages. However, we see that the model predicts an even slower decay than observed in these conditional probabilities. In other words, at longer leads, persistence is somewhat too high in the model. When averaged over all ages, this yields the probabilities in Table 2 being higher than their empirical counterparts.

To see this more clearly, consider Figure 12. This figure compares model with data along the dimension of the probability (weighted by the age distribution) of distress at various leads beyond an initial distress period. As before, we see that while the model produces less persistence at extremely short leads, it predicts slightly greater persistence at longer horizons. As we will show, despite the imperfect fit here, it is far better than a wide class of alternatives that might initially be seen as candidates.

An aspect of financial distress that sheds light on the source of its concentration across

Figure 10: The Concentration of Financial Distress: FRBNY/Equifax data vs Benchmark model





Figure 11: The Persistence of Financial Distress Over the Life Cycle: FRBNY/Equifax data (left) and Benchmark model (right)



households is the distribution of the number of spells of distress across households (for those who have experienced it). The left panel of Figure 13 shows that in the data, the bulk of people who experienced distress also experienced only one spell. The right panel of Figure 13 suggests the model has a similar implication, though it understates the share of individuals with only one spell.

The preceding intuition can be sharpened by referring to Figure 14. This displays the fraction of an individual's life spent in financial distress. As is clear the data place considerable mass on outcomes involving large proportions of life spent in distress. Indeed, about 30% of individuals spend a quarter or more of their lives in financial distress. As seen in the right-hand panel, the model does fairly well in accounting for this, though underpredicting somewhat the proportion of those spending very small proportions of their lives in distress.

Figure 12: Average Persistence of Financial Distress: FRBNY/Equifax data (left) and Benchmark model (right)




Figure 13: The Duration of Financial Distress (spells): FRBNY/Equifax data (left) and Benchmark model (right)





Figure 14: The Duration of Financial Distress (share of life): FRBNY/Equifax data (left) and Benchmark model (right)



3.2 Why are Informal Default and Discount-factor Heterogeneity Essential?

Having demonstrated that a relatively simple extension of the canonical approach of Livshits et al. (2007) is sufficient to allow that workhorse model to account well for the dynamics of financial distress as a nontargeted outcome, we now detail results that suggest that these extensions are also likely to be necessary. In this sense, the facts we establish on financial distress suggest that borrowers may differ very fundamentally from each other—in ways simply not captured by income risk or by the precise specification of options for debt default.

3.2.1 Distress as Bankruptcy, RIP Income Risk

We begin by focusing on a simpler model that coincides very closely with that of Livshits et al. (2007). In this setting, default is modeled as bankruptcy (BK), and income is, as in that benchmark, of the standard “RIP” variety. Critically, the model is one in which default removes debt obligations altogether and imposes costs in the current period. Table 4 presents the parameters that are externally chosen for this model. Many are the same as in our benchmark model, with the exception of the roll-over interest rate on delinquent debt and the filing cost of bankruptcy, which are both mechanically set to zero given the structure of this model. Given the simplicity of this model, the parameters that remain to be pinned down are the discount factor β (which is common across all individuals) and the income garnishment threshold when in bankruptcy τ . For ease of comparison with the benchmark model, this model is calibrated to match similar facts on bankruptcy, as seen in Table 5 below, while the implied parameter values are displayed in Table 6. The reader is referred to the Appendix for a detailed statement of the environment.

Table 4: Bankruptcy Model Parameters Determined Externally

Parameter	Value
σ , Coefficient of relative risk aversion	2.000
r , Risk free interest rate	3.000%
W , Retirement age	65
A_0 , Replacement ratio	0.71
A_1 , Replacement ratio	-0.045
A_2 , Replacement ratio	0.14
RIP process	
σ_ε^2 , Variance of permanent shocks	0.05
σ_e^2 , Variance of transitory shocks	0.01
HIP process	
$\{\sigma_\alpha^2, \sigma_\beta^2, \sigma_\varepsilon^2, \sigma_e^2\}$	$\{0.022, 0.00038, 0.047, 0.029\}$
ρ_z	0.821
$corr_{\alpha\beta}$	-0.23

Table 5: Bankruptcy Model Calibrations

Moment	Data	BK RIP	BK HIP
FD rate at age 30	14.24	14.95	14.17
FD rate at age 50	9.30	6.53	9.30
Pop. that never defaults	64.23	44.00	35.19
FD share of top 20% of pop.	88.07	59.02	54.0
Avg. cond. FD rate at 2 years	0.520	0.261	0.275
Avg. cond. FD rate at 4 years	0.340	0.152	0.199
Avg. cond. FD rate at 6 years	0.231	0.116	0.167
Avg. cond. FD rate at 8 years	0.171	0.099	0.145
Avg. cond. FD rate at 10 years	0.153	0.089	0.133


Note: Numbers in bold denote targeted moments in the calibration of each model.

Table 6: Calibrated Parameters for Bankruptcy Models

Parameter	BK RIP	BK HIP
Discount factor β	0.839	0.687
Earnings threshold in default τ	4.276	3.294

The results are fairly self-explanatory. Not only does this model fail to generate the proportion of those who never experience financial distress, those who do remain in it for far too short a time. One way to see this immediately is to look at Figure 15.

Figure 15: The Persistence of Financial Distress Over the Life Cycle: FRBNY/Equifax data (left) and BK RIP model (right)



Similarly, we can see from the bar chart in Figure 16 that while by construction this model misses on persistence at the one year mark (since repeat bankruptcy is not allowed), its implications for persistence are counterfactual beyond that.

Lastly, but importantly, we see in Figure 17 that this simpler alternative fails substantially at generating the concentration of financial distress.

Figure 16: Average Persistence of Financial Distress: FRBNY/Equifax data (left) and BK RIP model (right)





Figure 17: The Concentration of Financial Distress: FRBNY/Equifax data vs BK RIP model



3.2.2 Distress as Bankruptcy, HIP Income Risk


The inability of the preceding model to capture the facts need not lead one to the much richer setting that we ultimately defined as the baseline model. Indeed, a natural first step before broadening the notion of financial distress, and certainly before allowing for the discount factor heterogeneity that we will eventually allow for is to ask: Have we represented income risk well? Specifically, it is relevant to gauge if the model's failure is driven by an income persistence that is a poor representation of variation in income prospects in the cross-section (within the bankruptcy-only model). To do this, we now extend the previous model to allow for heterogeneity in earnings profiles, i.e. the so-called HIP specification of Guvenen (2009). This modification imparts permanent differences in borrowers' lifetime income growth rates. Thus, log income now is of the form: $y_{i,n} = l_n + \alpha_i + \beta_i n + z_{i,n} + e_{i,n}$. As before, l_n denotes the life cycle component common to all households of age n and $e_{i,n}$ is a transitory

component. The term $\alpha_i + \beta_i n$ is the life cycle component that is household-specific. Next, $z_{i,n}$ follows an AR(1) process: $z_{i,n} = \rho_z z_{i,n-1} + \varepsilon_{i,n}$. As in Guvenen (2009), we assume the random vector (α_i, β_i) is distributed across households with zero mean, variances of σ_α^2 and σ_β^2 , and correlation of $\text{corr}_{\alpha\beta}$. Finally, $\varepsilon_{i,n}$, and $e_{i,n}$ are normally distributed with variances σ_ε^2 , and σ_e^2 , respectively. We calibrate this model in the same fashion as the RIP version. The third column of Table 5 displays the calibration of this model and some basic nontargeted moments. The second column of Table 6 displays the implied parameter values consistent with these empirical targets. As with the previous models, this model matches several basic cross-sectional observations on financial distress, but does less well on nontargeted moments than the baseline model.

As seen in Figure 18, the baseline model with the HIP process generates more persistence, but only later in life. Figure 19 shows that again, by construction, this model misses on the persistence at the one year mark, but also severely underestimates the persistence several years out. Lastly, and just as under RIP income risk, Figure 20 shows the share of individuals who never default over 15 years is counterfactually low in this model.

The stark inability of either variant of the benchmark consumer default model to generate the proportion of nondefaulters is a key observation, and will inform the development of extensions that do better. The reason is this: Any model that matches the unconditional incidence of financial distress over the life cycle but fails on persistence will *necessarily*—and grossly, in the cases above—overstate the probability that an individual experiences financial distress. After all, 65% of the agents in the model will hit distress, compared with just 35% in the data! On the other hand, this observation clarifies that capturing persistence is absolutely central to providing a reasonable description of people’s experience over the life cycle as they relate to financial distress.

Figure 18: The Persistence of Financial Distress Over the Life Cycle: FRBNY/Equifax data (left) and BK HIP model (right)



So far, we have presented results that show that under the two major categories of income risk considered in the literature, financial distress is predicted to be far more transitory than in the data. The forces underlying this turn out to be the ones at work in general life cycle consumption/savings models. Specifically, in the model above, under either income process, life cycle consumption leads households, all else equal, to borrow. Income risk, however,

Figure 19: Average Persistence of Financial Distress: FRBNY/Equifax data (left) and BK HIP model (right).





Figure 20: The Concentration of Financial Distress: FRBNY/Equifax data and BK HIP model.



pushes consumers away from borrowing (as long as any precautionary motive is present, as it is in our parametrization of preferences). As long understood, in models without default, such as the classic work of Gourinchas and Parker (2002), the net result is for consumption to track income early in the life cycle—just as in the data and for it to then fall below it as retirement nears—just as in the data. The question is then, how does the ability to default alter this basic finding? Athreya (2008) finds that for plausible (RIP) specifications of income risk, the ability to default, when quantitatively disciplined, restricts credit access substantially. This results in the model behaving similarly to one in which individuals are simply more constrained. (In other words, the state-contingency provided by the ability to obtain debt relief is not very valuable.). This behavior rules out very large defaultable debts from building up frequently in early life and virtually eliminates debts altogether as agents

near retirement. By contrast, the data clearly show that the incidence of financial distress does not vanish even very late in (pre-retirement) life.

The previous two models, taken together, suggest that the particulars of individual income risk, while of course important in ways established by the work cited above, is not at the heart of explaining the persistence of financial distress. This is, a priori, not an obvious conclusion. After all, consumption smoothing motives have long been understood to hinge on the particulars for income risk. Most notably, a setting (with infinite-lived agents, especially) where income shocks are very persistent is one in which borrowing to smooth consumption makes less sense than simply lowering consumption to reflect a new, and substantially worse, reality. Conversely, the more purely transitory the income risk, the more that individuals will incur debts to maintain smooth consumption. And yet, we have seen that fairly disparate representations of income risk, once models are calibrated to match basic cross-sectional observations, fail completely to describe the persistent state of affairs that characterizes financial distress in US data.

3.2.3 Distress as Delinquency and Bankruptcy

However, neither of the models just considered distinguish between nonpayment of debt and bankruptcy, while our measure of distress in the data is that of severe delinquency. It is therefore important to assess whether this omission is important in the failure of the models to match the persistence (and concentration) of financial distress. A main reason, at least up front, to worry about this omission is that bankruptcy carries fixed costs arising from income loss or other fees. Thus, one would not expect it to be used repeatedly. By contrast, financial distress as delinquency does not impose such costs. A model that allows for such an option therefore seems warranted.

The model we now consider follows that of Athreya et al. (2017) and allows for informal delinquency where, as in the baseline model, borrowers may opt to cease repayment and pay a “penalty” interest rate on existing debt. This model (DQBK) is disciplined by three facts: The rate of financial distress at age 30, at age 50, and the formal bankruptcy rate. The Appendix presents this model’s performance in matching these facts and its calibrated parameters.

Focusing on the empirical performance of this model, from Figure 21 we see that this model matches the concentration of financial distress quite well, albeit missing slightly concentration of the top 20% of the population. As for persistence across the life cycle, we see from Figure 22 that while persistence is reasonably captured during middle age, it is far away from the data at both young and old ages. The miss at old ages highlights the importance of discount factor heterogeneity, whereby effectively impatient older individuals still find themselves systematically in distress (note that Figure 11 displays no such decline in persistence).

Figure 21: The Concentration of Financial Distress: FRBNY/Equifax data and DQBK model




Figure 22: The Persistence of Financial Distress Over the Life Cycle: FRBNY/Equifax data (left) and DQBK model (right)





Figure 23: Average Persistence of Financial Distress: FRBNY/Equifax data (left) and DQBK model (right)



The observant reader will have noticed that this model (DQBK) seems to match many of the facts as well, or nearly as well, as our benchmark model, which features substantial heterogeneity in discounting. But the comparisons thus far provide an incomplete picture. To see this, in Table 8 of the Appendix we show that the best calibration of the DQBK model implies a discount rate for all agents of roughly 17% (i.e. $\beta=0.83$). This is implausible for two main reasons. First, and most importantly, this value is at odds with those used in the vast majority of the vast literature on consumption and savings. Indeed, in that work, the modal discount factor has been equal, or close, to 0.95. Even models like ours that aim to account for default, and hence may be driven to allow lower values the discount factor, do not imply values anywhere near as low as 0.83. For example, the benchmark model of Livshits et al. (2007) the discount factor uses a discount factor of 0.94. It is, of course though not random why baseline models with a single discount factor generally impose discount factors much larger than in the DQBK model: Essentially any model that aims to match wealth accumulation over the life cycle will fail completely at lower values. Figure 24 compares the performance of both the DQBK and Benchmark models relative to data from the 1998 Survey of Consumer Finances (SCF).¹³ This figure makes clear just how drastically the former underpredicts mean wealth accumulation (expressed relative to income) over the life cycle relative to both the data and the Benchmark model. While not shown here, uniformly low discount factors are similarly problematic for the “BK” model that defined financial distress as bankruptcy, illustrating yet another shortcoming of that model for understanding financial distress.¹⁴ While failures of this magnitude are perhaps inconsequential in settings not trying to understand balance sheet positions, they are unlikely to be so for addressing the question of interest here. Second, as documented at the outset, analysis of a wide variety of empirical evidence (most recently that of Fulford and Schuh (2017) and Parker (2017)) provides strong support for the view that the null hypothesis should be in favor of models that allow heterogeneity in discounting rather than not.


3.2.4 Why don’t models that disallow variation across agents in discount rates capture persistence well?

The answer lies in the equilibrium incentives to use any form of default. In smoothing consumption over both the life cycle and across contingencies agents in the model are, as usual, seeking to equalize marginal utilities along both dimensions. In the presence of the option to default, the marginal cost of moving consumption across dates and states varies with the agents current state. Whenever income is persistent, current income helps predict future income and hence default incentives. As a result those faring poorly at present face high marginal costs of credit all else equal. This in turn makes borrowing unattractive all else equal. Nonetheless, all models are calibrated to yield plausible overall default rates. This, in turn, demands that borrowers be willing to pay high costs for credit, and is why in the case with no discount factor heterogeneity, the single discount rate produced in the calibration ($\beta=0.83$) is far higher than implied by environments that abstract from consumer debt, and especially ones that abstract from consumer default. A common, but high, discount rate, however runs into difficulty because it almost immediately implies a widespread willingness

¹³We chose this cross-section as it most closely aligns with the 1999 cohort from our Equifax sample.

¹⁴Results are available upon request

Figure 24: Wealth to income ratios: SCF 1998 data, Benchmark model, and DQBK model



to borrow and default, all else equal. Yet the data on the concentration of distress suggests almost the opposite: a few borrowers account for most of distress seen. At the same time, the presence of costs of default make this option not one to be entered into repeatedly. In the case of formal bankruptcy, there are fixed costs, and in the case of delinquency, borrowers need to be induced into paying the penalty rate (20%) repeatedly if the model is to predict correctly the persistence of financial distress. Neither of those is likely under a baseline setting in which there is a single agent-type. By contrast, once this assumption is relaxed, we can account both for the persistence (at most ages over the life cycle) and concentration of financial distress. In particular, the allowance for heterogeneity in discounting yields the intuitive change (relative to the baseline) of a two tier outcome whereby a majority (56%) are characterized as being effectively impatient ($\beta=0.74$) while a minority (44%) are extremely patient ($\beta=1.00$). In sum, the data on the persistence of financial distress and its concentration within a small group of borrowers make it essentially necessary (and sufficient too) to allow for discounting to vary in the population.

4 Conclusion

This paper establishes first that using recently available proprietary panel data, while many (35%) US consumers experience financial distress at some point in the life cycle, most of financial distress events are primarily accounted by a much smaller proportion of consumers in persistent trouble. For example, about 10% are distressed for more than a quarter of the life cycle, and less than 10% of borrowers account for half of all distress. Second, we show that these facts can be largely accounted for in a straightforward extension of a workhorse model of defaultable debt that accommodates a simple form of heterogeneity in time preference but not otherwise. Specifically, the data are strongly consistent with the presence of a subset of effectively impatient consumers, rendered so potentially by a host of additional factors not

modeled here. Future work that allows for more detail on household-level economic dynamics is therefore essential to more deeply understand the sources of this apparent heterogeneity, certainly before reaching any conclusions that “implicate” individuals in their fates simply via differences in their preferences.

References

- AMPUDIA, M., H. VAN VLOKHOVEN AND Z. DAWID, “Financial Fragility of Euro Area Households,” *Journal of Financial Stability* 27 (2016), 250–262.
- ATHREYA, K., “Default, Insurance, and Debt over the Life-Cycle,” *Journal of Monetary Economics* 55 (May 2008), 752–774.
- ATHREYA, K., J. M. SÁNCHEZ, X. S. TAM AND E. R. YOUNG, “Bankruptcy and delinquency in a model of unsecured debt,” Technical Report, Federal Reserve Bank of St. Louis, 2017, forthcoming *International Economic Review*.
- BECKER, S. G. AND C. B. MULLIGAN, “The Endogenous Determination of Time Preference,” *The Quarterly Journal of Economics* 112 (August 1997), 729–758.
- BELZIL, C. AND J. HANSEN, “Subjective Discount Rates, Intergenerational Transfers and the Return to Schooling,” Technical Report, IZA Discussion paper series, No. 60, October 1999.
- BLUNDELL, R., L. PISTAFERRI AND I. PRESTON, “Consumption Inequality and Partial Insurance,” *American Economic Review* 98 (December 2008), 1887–1921.
- BRUNETTI, M., G. ELENA AND T. COSTANZA, “Is Financial Fragility a Matter of Illiquidity? An Appraisal for Italian Households,” *Review of Income and Wealth* 62 (December 2016), 628–649.
- CAGETTI, M., “Wealth Accumulation Over the Life Cycle and Precautionary Savings,” *Journal of Business and Economic Statistics* 28 (2003), 339–353.
- CHATTERJEE, S., D. CORBAE, M. NAKAJIMA AND J.-V. RÍOS-RULL, “A Quantitative Theory of Unsecured Consumer Credit with Risk of Default,” *Econometrica* 75 (June 2007), 1525–1591.
- CHATTERJEE, S., D. CORBAE AND J. V. RÍOS-RULL, “A Finite-Life Private-Information Theory of Unsecured Consumer Debt,” *Journal of Economic Theory* 142 (2008), 149–177.
- FREDERICK, S., G. LOEWENSTEIN AND T. O’DONOGHUE, “Time Discounting and Time Preference: A Critical Review,” *Journal of Economic Literature* XL (June 2002), 351–401.
- FULFORD, S. L. AND S. SCHUH, “Credit Card Utilization and Consumption over the Life Cycle and Business Cycle,” Technical Report, Federal Reserve Bank of Boston, September 2017.

- GORBACHEV, O. AND M. J. LUENGO-PRADO, "The Credit Card Debt Puzzle: The Role of Preferences, Credit Risk, and Financial Literacy," Technical Report, Federal Reserve Bank of Boston, 2016, research Department Working Paper 16-6.
- GOURINCHAS, P.-O. AND J. A. PARKER, "Consumption Over the Life Cycle," *Econometrica* 70 (January 2002), 47–89.
- GROSS, D. B. AND N. S. SOULELES, "Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data.," *Quarterly Journal of Economics* 117 (February 2002), 149–185.
- GUVENEN, F., "An Empirical Investigation of Labor Income Processes," *Review of Economic Dynamics* 12 (2009), 58–79.
- HALL, E., ROBERT AND F. MISHKIN, "The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households," *Econometrica* 50 (March 1982), 461–481.
- HATCHONDO, J. C., L. MARTINEZ AND J. M. SÁNCHEZ, "Mortgage Defaults," *Journal of Monetary Economics* 76 (November 2015), 173–190.
- HAUSMAN, J. A., "Individual Discount Rates and the Purchase and Utilization of Energy-Using Durables," *Bell Journal of Economics* 10 (Spring 1979), 33–54.
- JAPPELLI, T. AND M. PAGANO, "The Welfare Effects of Liquidity Constraints," *Oxford Economic Papers* 51 (1999), 410–430.
- JAPPELLI, T., M. PAGANO AND M. DI MAGGIO, "Households? Indebtedness and Financial Fragility," *Journal of Financial Management Markets and Institutions* (2013), 26–35.
- KAPLAN, G. AND G. L. VIOLANTE, "How Much Consumption Insurance beyond Self-Insurance?," *American Economic Journal: Macroeconomics* 2 (October 2010), 53–87.
- LAWRANCE, E. C., "Poverty and the Rate of Time Preference: Evidence from Panel Data," *Journal of Political Economy* 99 (February 1991), 54–77.
- LIVSHITS, I., J. MACGEE AND M. TERTILT, "Consumer Bankruptcy: A Fresh Start," *The American Economic Review* 97 (March 2007), 402–418.
- LUSARDI, A., "Americans' Financial Capability," NBER Working Paper No. 17103, June 2011.
- LUSARDI, A., D. SCHNEIDER AND P. TUFANO, "Financially Fragile Households: Evidence and Implications," *Brookings Papers on Economic Activity* (Spring 2011), 83–134.
- MEIER, S. AND C. SPRENGER, "Present-Biased Preferences and Credit Card Borrowing," Technical Report, Federal Reserve Bank of Boston, 2017, research Department Working Paper 07-3. Forthcoming *American Economic Journal: Applied Economics*.

- MUSTRE-DEL RÍO, J., “Wealth and Labor Supply Heterogeneity,” *Review of Economic Dynamics* 18 (2015), 619–634.
- PARKER, J. A., “Why Don’t Households Smooth Consumption? Evidence from a \$25 Million Experiment,” *American Economic Journal: Macroeconomics* 9 (2017), 153–183.
- SAMWICK, A. A., “Discount Rate Heterogeneity and Social Security Reform,” *Journal of Development Economics* 57 (October 1998), 117–146.
- WARNER, J. T. AND S. PLEETER, “The Personal Discount Rate: Evidence from Military Downsizing Programs,” *The American Economic Review* 91 (March 2001), 33–53.
- ZELDES, S., “Consumption and Liquidity Constraints: An Empirical Investigation,” *Journal of Political Economy* 97 (April 1989), 305–346.

A Data and Moment Construction

This appendix provides a description of the data used. All our empirical work leverages information from Federal Reserve Bank of New York Consumer Credit Panel/Equifax, unless otherwise noted. We trimmed our sample such that individuals missing in any quarter from 1999Q1 to 2017Q2 are dropped. Additionally, we restrict attention to individuals who enter the sample in 1999Q1 between the ages of 25 and 55.

Unconditional Fraction of Individuals in DQ. Unconditional fraction of individuals in delinquency (DQ), also called unconditional probability of being in DQ, is calculated by finding the ratio of DQ debt to total number of individuals. DQ debt is computed as the sum of balances of all delinquent accounts if an individual is more than 120 DPD, or Severe Derogatory, i.e., $\text{DQ_debt}_{i,j} = \text{crtr_attr111} + \text{crtr_attr112}$, for individual i at age j . A dummy variable $\mathbb{1}_{DQ_{i,j}}$ is defined for all individuals, where $\mathbb{1}_{DQ_{i,j}} = 1$ if $\text{DQ_debt}_{i,j} > 0$. Note that if individual is in delinquency at least one quarter on a particular age, $\mathbb{1}_{DQ_{i,j}} = 1$. Then unconditional fraction of individuals in DQ is calculated as

$$\text{Unconditional fraction of individuals in DQ} = \frac{\sum_{i=1}^{N_j} \mathbb{1}_{DQ_{i,j}}}{N_j}$$

Unconditional Fraction of Debt in DQ. Similarly, unconditional fraction of debt in DQ is computed by finding the ratio of DQ debt to total debt. Total debt is computed as the sum of balances of all accounts, i.e., $\text{Total_debt}_{i,j} = \text{crtr_attr107} + \text{crtr_attr108} + \text{crtr_attr109} + \text{crtr_attr110} + \text{crtr_attr111} + \text{crtr_attr112}$. Then unconditional fraction of debt in DQ is

$$\text{Unconditional fraction of debt in DQ} = \frac{\sum_{i=1}^{N_j} \frac{\text{DQ_debt}_{i,j}}{\text{Total_debt}_{i,j}}}{N_j}$$

Conditional Probability of Being in DQ. We compute the probability of being in DQ conditional on being in DQ h years ago as

$$\text{Conditional probability of being in DQ debt} = \frac{\sum_{i=1}^{N_j} \mathbb{1}_{DQ_{i,j}} \cdot \mathbb{1}_{DQ_{i,j+h}}}{\sum_{i=1}^{N_j} \mathbb{1}_{DQ_{i,j}}}$$

It is important to note that $\mathbb{1}_{DQ_{i,j}}$ does not contain everyone who is in delinquency at age j when computing conditional probability. In fact, any individual whose age is $j + h > j^*$ where j^* is maximum age in the sample period is dropped. For example, if an individual i is in delinquency at age 40 in year 2014, $j^* = 43$, and this individual is excluded from $\mathbb{1}_{DQ_{i,40}}$ in the computation for conditional probability for age greater than 43, since we do not have data beyond 2017. This individual is not excluded when computing unconditional probability.

Unconditional Probability of Reaching Credit Limit. Unconditional probability of reaching credit limit is calculated by finding the ratio of individuals reaching credit limit to total number of individuals. Another dummy variable $\mathbb{1}_{Credit_{i,j}}$ is defined for all individuals, where $\mathbb{1}_{Credit_{i,j}} = 1$ if bank balance $\geq$ credit limit, i.e., $crtr_attr169 \geq crtr_attr180$. Similarly, if the individual has reached credit limit at least one quarter on a particular age, $\mathbb{1}_{Credit_{i,j}} = 1$. Then unconditional probability of reaching credit limit is calculated as

$$\text{Unconditional probability of reaching credit limit} = \frac{\sum_{i=1}^{N_j} \mathbb{1}_{Credit_{i,j}}}{N_j}$$

Conditional Probability of Reaching Credit Limit. Similarly, probability of reaching credit limit, conditional on reaching credit limit h years ago, is computed as

$$\text{Conditional probability of reaching credit limit} = \frac{\sum_{i=1}^{N_j} \mathbb{1}_{Credit_{i,j}} \cdot \mathbb{1}_{Credit_{i,j+h}}}{\sum_{i=1}^{N_j} \mathbb{1}_{Credit_{i,j}}}$$

Average Life of DQ. Average life of DQ for individual i is computed as the ratio of total number of quarters i is in DQ ($\mathbb{1}_{DQ_{i,j}} = 1$) to the total number of quarters in the sample period for i . Let $DQnum_i$ be the total number of quarters i is in DQ, and let T_i denote the total number of quarters in the sample period for i . Then

$$\text{Average life in DQ for } i = \frac{DQnum_i}{T_i}$$

Note that Figure 4 excludes individuals who do not spend any quarter in DQ because the large proportion of the population that does not enter DQ distorts the scale of the histogram.

Delinquency Spell Number. A delinquency spell begins when the individual is in DQ ($\mathbb{1}_{DQ_{i,j}} = 1$) in the current quarter but was not in DQ the preceding quarter. Similarly, a delinquency spell ends when the individual is not in DQ in the current quarter but was in DQ preceding quarter. If the first and last observation is in DQ, we take that quarter to be the start or end of the DQ respectively. Note that an individual can have multiple delinquency spell throughout his life. Also note that x-axis of Figure 5 has been trimmed to 10 for illustrative purpose. The original scale spans to 14, but the cumulative density between 11 to 14 spells accounts for less than 0.1%.

Lorenz Curves. Lorenz curves are calculated using two measures: being in DQ, and reaching credit limit. After sorting out the individuals in a nondecreasing order by $DQnum_i$, share of DQ (y-axis of Lorenz curve) is computed as the following

$$\text{Share of DQ for } \hat{i} = \frac{\sum_{i=1}^{\hat{i}} DQnum_i}{\sum_{i=1}^N DQnum_i}$$

Share of DQ for $\hat{i}$ is then plotted against the share of population that is given by $\frac{\hat{i}}{N}$. Similar computation applies for credit limit.

Delinquency Intensity. Delinquency intensity is computed as the average ratio of debt in DQ to total debt among people that have entered DQ. Hence it is

$$\text{Delinquency Intensity} = \frac{\sum_{i=1}^{N_j} \frac{\text{DQ_debt}_{i,j}}{\text{Total_debt}_{i,j}}}{\sum_{i=1}^{N_j} \mathbb{1}_{DQ_{i,j}}}$$

Alternative measure of delinquency intensity is calculated by taking number of bankcard 120 DPD or Severe Derogatory. Let $\text{Num_card}_{i,j} = \text{crtr_attr17} + \text{crtr_attr38}$, while $\text{Total_card}_{i,j} = \text{crtr_attr33} + \text{crtr_attr34} + \text{crtr_attr35} + \text{crtr_attr36} + \text{crtr_attr37} + \text{crtr_attr38}$. Define $\mathbb{1}_{Card_{i,j}} = 1$ if $\text{Num_card}_{i,j} > 0$. Then it is computed as


$$\text{Delinquency Intensity} = \frac{\sum_{i=1}^{N_j} \frac{\text{Num_card}_{i,j}}{\text{Total_card}_{i,j}}}{\sum_{i=1}^{N_j} \mathbb{1}_{Card_{i,j}}}$$

Delinquent Debt Figure 9 is computed by taking 50th, 75th and 90th percentile of $\text{DQ_debt}_{i,j}$ by age. Note that the amount of DQ debt has been inflation-adjusted to 2017 January dollars using seasonally adjusted CPI from US Bureau of Labor Statistics.

B Cross-state Comparisons


To ensure that our findings are not simply driven by the vagaries of any single state of the union in the data, Figures 25 and 26 present the life cycle incidence and persistence of financial distress across the six most populous states in the data. As is clear, not only are the qualitative patterns extremely similar across states, so are the quantities. Thus, we see that across the US, financial distress patterns are very similar, and this plausibly amenable to analysis within a model framework that abstracts from what might have seemed, a priori, relevant differences across states.

Figure 25: The Persistence of Financial Distress Over the Life Cycle and Across States (debt)



Source: Federal Reserve Bank of New York Consumer Credit Panel/Equifax

Figure 26: The Persistence of Financial Distress Over the Life Cycle and Across States (credit limit)



Source: Federal Reserve Bank of New York Consumer Credit Panel/Equifax

C A Model of Distress as Bankruptcy

In this section we provide details of the model of distress as bankruptcy from Section 3.2.1 in the main text.

We assume that each period households may default on existing debt. Like in our benchmark model in the main text, households trade-off the advantages and disadvantages of bankruptcy. The key advantage is the discharge of debts: current period expense obligations are eliminated and in the period after bankruptcy, debt is set at zero. Thus, a household with too much debt may find it beneficial to file for bankruptcy. There are two disadvantages of doing so, however. In the period of bankruptcy, a proportion of income, τ , is lost.¹⁵ Additionally, in that period, consumption equals income—neither saving nor borrowing are allowed. In this environment, lifetime utility can be written as

$$G_{i,n}(z, \varepsilon, a) = \max\{\underbrace{V_{i,n}(z, \varepsilon, a)}_{\text{Pay}}, \underbrace{B_{i,n}(z, \varepsilon)}_{\text{Bankruptcy}}\} \quad (6)$$

where V and B (defined below) are lifetime utilities for households paying back the debt and filing bankruptcy, respectively. This means that a household has the choice of filing bankruptcy. The policy function R indicates whether the household pays back the debt (repay) or not,

$$R_{i,n}(z, \varepsilon, a) = \begin{cases} 1 & \text{if } V_{i,n}(z, \varepsilon, a) \geq B_{i,n}(z, \varepsilon), \\ 0 & \text{otherwise.} \end{cases}$$

Suppose the receives the opportunity to file for bankruptcy and chooses to do so. Then, lifetime utility is

$$B_{i,n}(z, \varepsilon) = u(\min\{y_{i,n}(z, \varepsilon), \tau\}) + \varrho_n \beta \mathbb{E}[G_{i,n+1}(z', \varepsilon', 0)|z]. \quad (7)$$

During the bankruptcy period, the household's consumption equals earned income up to a threshold $\tau > 0$. In the period after bankruptcy, the household will have no debt.

Now suppose the household pays back its debt. Then it faces the debt price $q_n(z, a')$ and lifetime utility

$$\begin{aligned} V_{i,n}(z, \varepsilon, a) &= \max_{\{a', c\}} u(c) + \varrho_n \beta \mathbb{E}[G_{i,n+1}(z', \varepsilon', a')|z], \\ \text{subject to} & \\ c + a' q_{i,n}(z, a') &= a + y_{i,n}(z, \varepsilon), \\ c &\geq 0. \end{aligned} \quad (8)$$

Equilibrium prices must imply zero-expected profits. In general, a price function $q_{i,n}(z, a')$ implies zero profits if the following equation is satisfied.

$$q_{i,n}(z, a') = \frac{1}{1+r} \varrho_n \mathbb{E}[R_{i,n}(z', \varepsilon', a')|z]. \quad (9)$$

¹⁵Chatterjee et al. (2008) build a model where no punishment is required after filing bankruptcy. There, asymmetric information is crucial to create incentives for debt repayment, because households signal their type by paying back their debt.

Looking at this equation it is very clear why the price function (or interest rates) depends on (a', z) . It depends on a' because it affects the bankruptcy decision, R , in each possible state. It depends on z because it determines the transition probability to each z' and therefore next period's earned income, y .

An equilibrium in this economy is a set of value functions, optimal decision rules for the consumer, default probabilities, and bond prices, such that equations (6) to (8) are satisfied and prices satisfy the zero-profit condition (9).

D Calibration of Model with Distress as Delinquency and Bankruptcy

In this section we present the model implied moments and calibrated parameters for the model presented in Section 3.2.3. The targeted moments are presented in Table 8, while the implied parameter values are presented in Table 8.

Table 7: DQBK Model Calibration

Moment	Data	Model
FD rate at age 30 (in %)	15.16	15.16
FD rate at age 50 (in %)	10.06	5.73
Average BK rate ages 25-65 (in %)	0.74	0.74
Pop. that never defaults (in %)	64.23	64.50
FD share of top 20% of pop. (in %)	88.07	81.84
Avg. cond. FD rate at 2 years	0.52	0.33
Avg. cond. FD rate at 4 years	0.34	0.23
Avg. cond. FD rate at 6 years	0.23	0.19
Avg. cond. FD rate at 8 years	0.17	0.16
Avg. cond. FD rate at 10 years	0.15	0.15

Note: Numbers in bold denote targeted moments in the calibration of the model.

Table 8: Calibrated Parameters for DQBK model

Parameter	Value
Discount factor β	0.83
Earnings threshold in DQ τ	3.85
Discharge shock to DQ debt γ	0.34