The Persistence of Financial Distress

Authors	Kartik Athreya, José Mustre-del-Río, and Juan M. Sánchez
Working Paper Number	2017-038C
Revision Date	July 2018
Citable Link	https://doi.org/10.20955/wp.2017.038
Suggested Citation	Athreya, K., Mustre-del-Rio, J., Sanchez, J.M., 2018; The Persistence of Financial Distress, Federal Reserve Bank of St. Louis Working Paper 2017-038. URL https://doi.org/10.20955/wp.2017.038

Published In	Review of Financial Studies
Publisher Link	https://doi.org/10.1093/rfs/hhz009

The Persistence of Financial Distress*

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

July 12, 2018

Abstract

Using proprietary panel data, we show that many US consumers experience financial distress (35% when distress is defined by severe delinquency, e.g.) at some point in their lives. However, most distress events are concentrated in a much smaller proportion of consumers in persistent trouble. While only 10% of consumers are distressed for more than a quarter of their lives, fewer 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 with informal default 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

*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 household financial distress (FD) in the United States, and to what extent can we understand them as arising from the choices of optimizing consumers who have access to credit and face uninsurable risks? 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 by estimating multiple state-of-the-art quantitative models of defaultable consumer debt over the life cycle.

The term "financial distress" can be defined in a variety of ways. Our primary definition is this: An individual will be said to be in financial distress in a given year if, in that year, at least one of their credit relationships (accounts) is at least 120 days past due, i.e., "severely delinquent." Because severe delinquency is, empirically, an expensive way to repeatedly roll-over debt, this definition plausibly captures financial distress. Put another way, because delinquency captures borrowers who face a high marginal cost of credit–it captures those with heavy debt and limited capacity to either self-insure or smooth consumption over time.

The preceding is not the only definition one might consider. In various places we therefore will also describe financial distress ways that emphasize the "intensive margin" —the volume of debt in delinquency, as well as one that emphasizes the "extensive margin" of credit-related problems. As for the latter, we will report a measure that answers the question "What is the proportion of consumers who have fully depleted the credit lines available to them?"

As we will show, the measures we consider all lead to a single, more general, conclusion: while many US consumers (35%, for our primary definition, e.g.) experience financial distress at some point in the life cycle, most distress events are primarily accounted for by a much smaller proportion of consumers in persistent trouble. In particular, distress incidence is nearly double its unconditional rate even a decade after the initial distress event. And while only about 10% of consumers are distressed for more than a quarter of the life cycle, just 10% of borrowers account for half of all distress. We also find that the persistence of FD is essentially invariant over the life cycle.[1]

The persistence of financial distress is important to measure and understand because it 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, with each episode

[1] In addition to these facts, and as with the facts on incidence described above, the persistence of financial distress is very similar across all 50 US states. The interested reader is referred to the Appendix.

not lasting long. 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. And the latter is what we find to be the case in the data. Our empirical findings therefore make clear that the risk of financial distress is one that is, in a sense, resolved early in life, with most borrowers knowing that they will face few problems with timely debt repayment in the years ahead, and a much smaller few knowing that they will face a future of repeated instances of severe delinquency.

Our work contributes in two ways. First, to our knowledge, our work is novel in providing a detailed description of the incidence, concentration, and dynamics of financial distress. Second, ours is the first to attempt to account for these facts. We will show that the facts of financial distress, along with overall wealth accumulation, can be largely accounted for through a straightforward extension of a workhorse model of defaultable debt with informal default that accommodates a simple form of heterogeneity in time preference, but not otherwise. In contrast, models without these features generate too little persistence and, as will show, fail to account even remotely for the empirics of life-cycle wealth accumulation.

By allowing for informal default, our model captures an empirically relevant pathway for (non) repayment, as reflected by the substantial delinquency rates observed in US data. By contrast, formal default (in its dominant "Chapter 7 Bankruptcy" form) is by construction very short-lived–it removes all unsecured debts–and thus fails to capture the ongoing difficulties experienced by households. In other words, informal default is the path for the many who are not ready to take the more extreme step of declaring bankruptcy but nonetheless face the difficulty–the financial distress–arising from potentially lengthy periods of costly-debt-rollover. To keep the model tractable for estimation, we follow Livshits et al. (2007), which was the first paper to allow for delinquency (in their work as an option in the period following a bankruptcy in response to "expense" shocks) by allowing for debt rollover at a "penalty" rate of interest.[2] And by allowing for discount-factor variation, we allow the model to generate the observed pattern of repeated and lengthy delinquency among a subset of the population, all while capturing overall life-cycle wealth accumulation patterns.

An additional motivation for our work, particularly our empirical efforts, is that in recent research the extreme events of bankruptcy or outright repudiation play an important role in helping

[2]See Athreya et al. (2015) for a richer model of informal default where, as in the model here, delinquency can be used at any time but where delinquent borrowers are potentially subject to optimal rate-resetting by incumbent lenders.

quantify the importance of limited commitment for allocations. Specifically, it is the observable rate of personal bankruptcy that provides a main target for the parameterization of the models. Recent work uses such models to analyze the implications of regulations (especially bankruptcy law) on outcomes. For example, recent reforms such as the Bankruptcy Abuse Prevention and Consumer Protection Act of 2005 (BAPCPA), and the effects of competing social insurance policies on credit use have been studied through versions of what is now a "standard default model" (e.g., Livshits et al. (2007), Chatterjee et al. (2007)). A consistent finding in this work (see also Athreya et al. (2009)) 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.[3] But as noted above, absent clear evidence that the baseline models used in these analyses capture well the time path of overall financial distress, there is reason for concern about the sensitivity of that finding.

Before proceeding, we stress that while our analysis suggests the presence of heterogeneity in discounting, 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.[4] 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.[5] Interest in the ability of the household to shield itself from susceptibility to shocks through the use of financial markets is,

[3] 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.

[4] Indeed, 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.

[5] http://www.cbsnews.com/news/the-financial-fragility-of-the-american-household/

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, within 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 (perhaps unsurprisingly) that those close to their limits increased borrowing by most, but (and more surprisingly) so did even those further away from their credit limit. A consensus might 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 one might broadly define to be those 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 imply too little persistence of distress. 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 a 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 findings also 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.[6] 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 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

[6] A 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 the early work of Lawrance (1991) 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.

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).[7] 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 provides the empirical analysis of financial distress in a proprietary panel data set (Equifax/NYFed Consumer Credit Panel). 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 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 the 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.[8] 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

[7] Relatedly, 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.

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

simply by behavior during and after the Great Recession. As stated earlier, we define an individual to be in financial distress in a given year if, in that period, they are recorded as having at least one severely delinquent (i.e., 120+ days past due) 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 solid black line shows the fraction of individuals in delinquency, not conditional on any credit-market status. What emerges is central to what follows: Financial distress, while relevant for consumers of all ages, is not widespread. The solid line in the figure begins near 10-14% among the young and falls below 10% later in life.

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

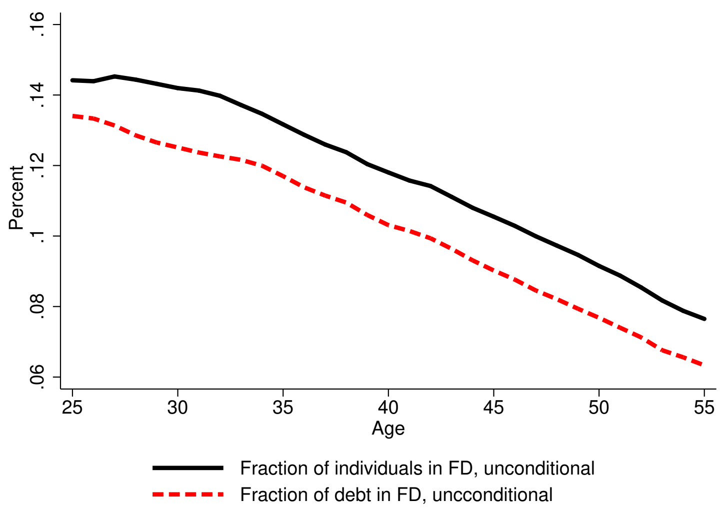

Source: See Appendix

As for the "intensive" margin of financial distress, consider instead the fraction of all debt at each age that is delinquent (again not conditional on a borrower's current or past credit status) represented by the dashed red line in Figure 1. We see that this measure follows a very similar pattern, with the youngest having the largest fraction of debt in delinquency (e.g., roughly 13% at age 25) and this proportion falling substantially to around 6% by age 55.

Next, to begin assessing the persistence of financial distress, it is natural to simply 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 financial distress is for US consumers. Conditional on being in distress today, the likelihood of being distressed in six years[9] (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 instead suggest the importance of heterogeneity in individual time preference.

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

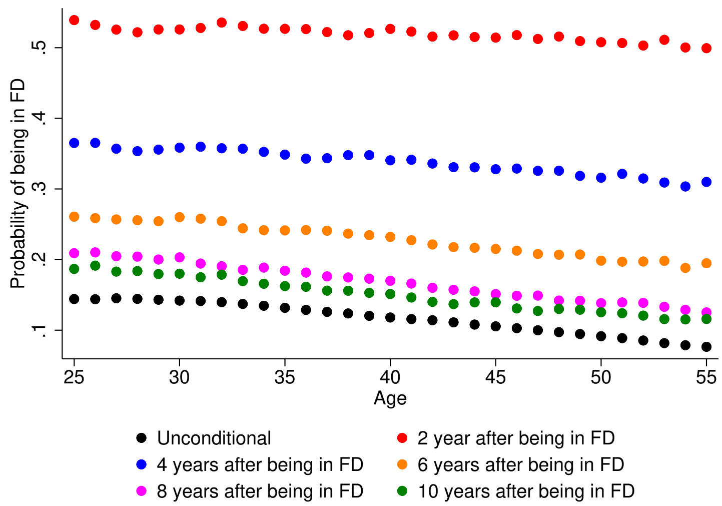

Source: See Appendix

As noted at the outset, 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 who have depleted their available credit (e.g, those who have "maxed out" their credit cards). To the extent that such credit, being unsecured, is expensive, the inability to arrange for more clearly represents at least a "fragility" or susceptibility to shocks, if not also direct distress. This metric is of the type seen in popular representations cited at the outset: that of a large subpopulation being unable to raise funds in an emergency. Figure 3 (which excludes those with 0 credit limit) shows that this notion of financial distress carries a very similar message: Limited borrowing capacity remains a prevalent issue for a small, but far from negligible, group of borrowers throughout the life cycle and, just as with severe-delinquency-based measures, displays substantial persistence.[10]

We now provide additional detail on the persistence of financial distress—defined, unless otherwise indicated, by our primary (severe-delinquency) definition. One point to keep in mind is that

[9]To be sure, note that this does not mean that the individual was necessarily in financial distress either continuously or at any one point during those six years.

[10]In 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)

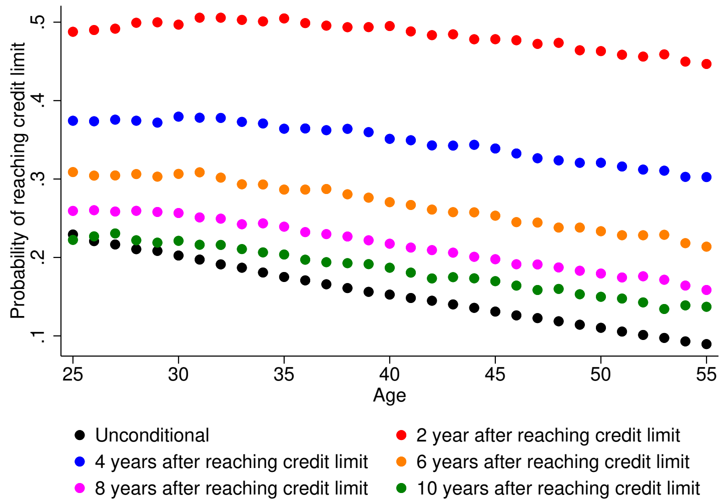

Source: See Appendix

the more transitory distress is, the less 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 on distress, this time measured by the distribution of time spent in financial distress. Specifically, we measure the proportion of the 18 years a consumer spends in our sample in 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 a 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!

Yet another way to gauge and 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)

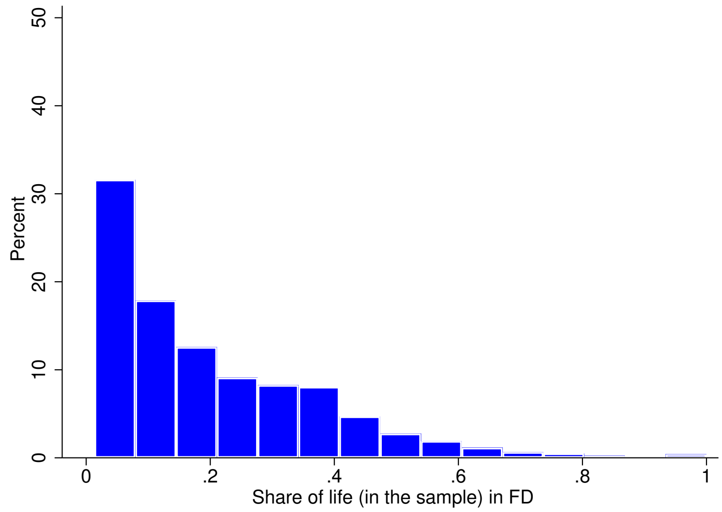

Source: See Appendix

Figure 5: The Duration of Financial Distress (spells)

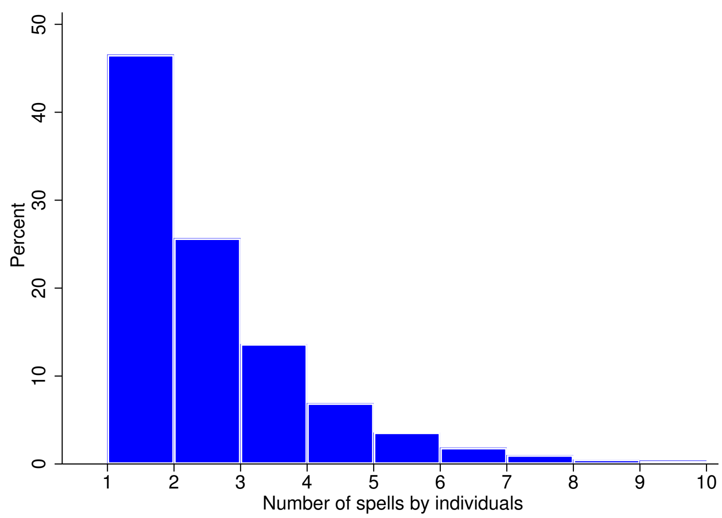

Source: See Appendix

Our measures imply that financial distress is "concentrated" in the population. Perhaps the most natural way to demonstrate this is via the Lorenz curves presented in Figure 6. They show that around 80% of financial distress is accounted for by less than 20% of people. This holds true whether we define financial distress as being severely delinquent (the solid black line) or having depleted available credit (dashed red line).

The measures presented thus far are primarily "extensive margin" measures: They are based

Figure 6: The Concentration of Financial Distress

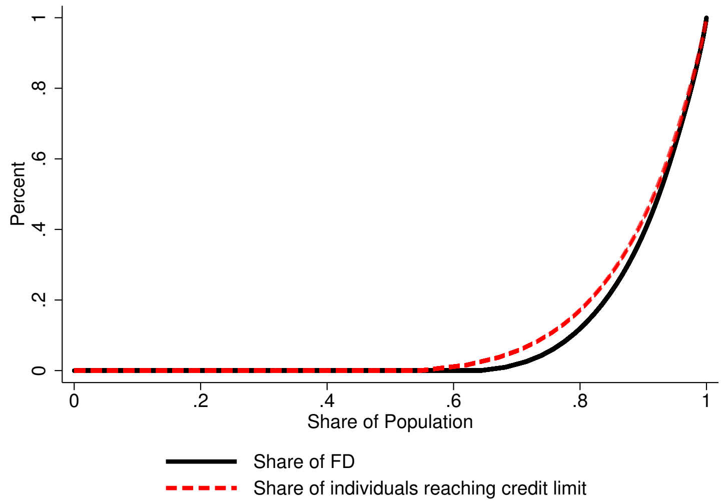

Source: See Appendix

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 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 on which 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 7 clearly indicates that financial distress among those facing it is "intense," measured in terms of either the average proportion across individuals of debt, or the average number of accounts severely delinquent. When debt is the measure (solid black line), we see that 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 by only 5 percentage points (88% early in the life cycle to 83% at older ages). A similarly flat life cycle profile is observed when intensity is measured as the average number of accounts severely delinquency (dashed red line). Overall, this figure suggests that when individuals are categorized as in financial distress based on our extensive margin measure, this is because most of their debt and most of their accounts are severely delinquent.

What about the size distribution of distressed debts? Figure 8 summarizes the amount of

Figure 7: The Intensity of Financial Distress Over the Life Cycle

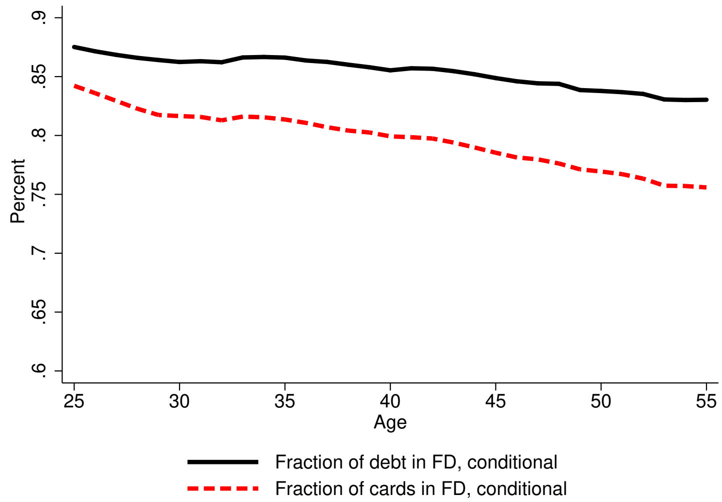

Source: See Appendix

delinquent debt by the 50th, 75th and 90th percentiles by age.[11] It shows that the median debt in delinquency is fairly stable over the life cycle, but the upper tail (measured by either the 75th or 90th percentile) 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.

Figure 8: Distressed Debt by Age

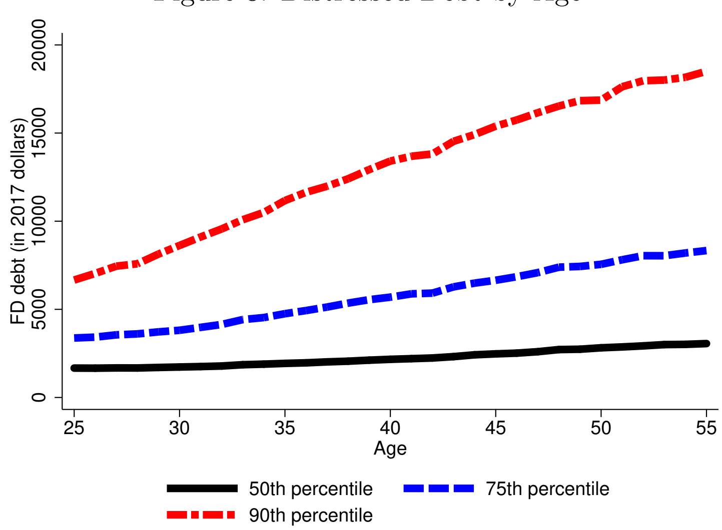

Source: See Appendix

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

Collecting the extensive- and intensive-margin empirics, we can summarize our findings as follows: Financial distress among US households, measured in a variety of ways, is driven by a relatively small proportion of individuals who experience significant and persistent debt repayment problems. And for those in it, financial distress is "intense" in the sense of applying to nearly all of their debts.

These facts suggest that financial distress may well be an important phenomenon, especially when from the point of view of a subset of individuals looking out, ex-ante, over a life cycle.[12] Discerning this importance, however, requires a model, which we turn to in the next section. We conclude this section by noting that any model that is successful in replicating these facts must also reproduce them in a context where debt repayment problems are significantly more common than the alternative (i.e., formal default via bankruptcy), and where agents, on average, nevertheless accumulate significant wealth over the life cycle.

3 Understanding Financial Distress

How well can the facts documented above, particularly the persistence of financial distress, be accounted for in a setting where households make empirically plausible choices over consumption, wealth, credit, and debt repayment? A natural starting place is the important work of Livshits et al. (2007) because it provides a benchmark life cycle consumption savings model in which debt may be repudiated formally or informally. Our model will feature two tractable extensions of this environment: Allowance for (i) the choice to informally default at any time and (ii) a simple form of heterogeneity in time preference. Using standard techniques, we estimate the key parameters of our benchmark model and show this extension is sufficient and necessary to account for the facts. Indeed, we show that when a fairly comprehensive set of alternative models are estimated to match the same set of facts, they fail to simultaneously generate the incidence, persistence, and concentration of financial distress–with a key sticking point being the ability to generate empirically reasonable wealth accumulation patterns over the life cycle.

[12] In 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.

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 or informal default. 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.[13] Moreover, in any period of delinquency, consumption equals income up to a threshold τ.[14] Second, and as 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—something that we let the data speak to in our estimation. 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

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

where V, B, and D are lifetime utilities for households paying back their debt, filing for bankruptcy, and being delinquent on their debt, respectively. Note that these functions are indexed by a a household's discount factor type j and age n. These functions take as arguments current household

[13] This 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.

[14] The remaining income is lost, for instance, when dealing with debt collectors.

wealth/debt and the household's income state. The latter is summarized by a permanent component z and a transitory component ε, both of which will be discussed in greater detail in the next section.

Next, the lifetime utility of bankruptcy is

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

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:

$$\begin{aligned} D_{j,n}(z,\varepsilon,a) \quad &= \quad u(\min\{y_n(z,\varepsilon),\tau\}) \\ &\quad +\varrho_n\beta_j\mathbb{E}\left[(1-\gamma)G_{j,n+1}(z',\varepsilon',(1+\eta)a) + \gamma G_{j,n+1}(z',\varepsilon',0)|z\right]. \end{aligned} \tag{3}$$

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

$$\begin{aligned} &V_{j,n}(z,\varepsilon,a) = \max\nolimits_{\{a',c\}} u(c) + \varrho_n\beta\mathbb{E}\left[G_{j,n+1}(z',\varepsilon',a')|z\right], \\ &\text{subject to} \\ &\qquad c + a'q_{j,n}(z,a') = a + y_n(z,\varepsilon), \\ &\qquad c \geq 0, \end{aligned} \tag{4}$$

where $q_{j,n}(z,a')$ 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:

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

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. Information is assumed complete: Lenders and borrowers are aware of all relevant state variables. Given this knowledge, lenders 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. Thus, we assume these variables are observable by lenders. Let the price of a debt issuance by a given borrower of type j,n be given as $q_{j,n}(z,a')$. 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:

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

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).

4 Calibration and Estimation

Our approach to model parameterization is standard: We assign parameter values in a two-step procedure.[15] First, we directly set values for a subset of the most "standard" parameters. Second, given these first-stage values, we formally estimate the remaining parameters and will then assess the model's performance in replicating the key empirics of financial distress. To estimate the most parsimonious models possible while still allowing for discount factor heterogeneity we make two assumptions. First, we fix the population shares of those with high- versus low-discount factor types to be equal, so $p_L = 0.5$.[16] Next, we assume that type H individuals have a discount factor of $\beta_H = 1$. These assumptions leave a total of three parameters (τ, β_L, γ) to be estimated in our benchmark model with discount factor heterogeneity (referred to as β-het), and informal delinquency and formal bankruptcy (referred to as DQBK).

Since preferences are unobservable, our allowance for preference heterogeneity should not be the only resort in accounting for the data on financial distress. Instead, for completeness, we will also consider both of the two best-known income process specifications: a restricted income profile (RIP) process and a heterogeneous income profile (HIP) process. The latter process incorporates ex-ante heterogeneity in income profiles. Thus, estimating our benchmark model with a HIP process helps assess how much unobservable ex-ante heterogeneity in preferences is needed above and beyond empirically observable ex-ante heterogeneity in income profiles. The former income process is often used because of its ability to match stylized facts about income over the life cycle in spite of its parsimonious structure. Thus, it serves as a useful benchmark compared to the literature.

Beyond two different income process specifications, we also consider two sequential deviations from our benchmark model. First, we drop the assumption of ex-ante heterogeneity in discount factors (referred to as No-het) but still allow for informal delinquency under both the RIP and HIP income processes. Importantly, because of the assumptions outlined above, these no-discount-factor-heterogeneity models still have the same number of parameters to be estimated. The relative fit of these No-het models compared to our benchmark models helps us assess the importance of

[15] See, e.g., De Nardi et al. (2016).

[16] This is a reasonable assumption for the following reasons. First, the facts on the concentration of financial distress mean that about 60 percent of individuals never experience a distress event in the 15 years we observe them. A lower bound approach might then lead to the assignment of all members of this group to the category of "patient" agents, and the remaining 40 percent would be impatient. However, this is likely to understate the measure of patient agents because of precautionary savings and fortunate realizations of idiosyncratic labor income draws over the life cycle, even some impatient types will never experience a distress event in the 15 years we observe them. Working under the (reasonable) presumption that the 60 percent of those who are never in distress almost certainly incorporates some impatient types, we therefore set the population shares to be equal.

discount factor heterogeneity. Second, we also drop the informal delinquency margin and consider a model with only bankruptcy (referred to No-DQ). Here too, we will still consider both RIP and HIP income processes. These No-DQ models have one extra degree of freedom, as bankruptcy guarantees full discharge of debt (this model is described in Appendix C). The relative fit of these No-DQ models compared to our benchmark or No-het models helps us assess the importance of modeling informal delinquency separate from formal bankruptcy.

4.1 Assigning First-Stage Parameters

Across all models, a period in the model refers to a year; households enter the model at age 25, retire at age 65, and die no later than at age 82. We set the risk free interest rate to 3% and assume households have constant relative risk aversion preferences over consumption setting $\sigma = 2$. In addition, we externally calibrate the parameters governing the income process, bankruptcy filing costs, retirement, and mortality. These are presented in Table 1. We also externally set the initial distribution of wealth-to-earnings to match the distribution of wealth-to-earnings of 25 year olds in the Survey of Consumer Finances between 1998 and 2016. Thus, initial financial conditions are constant across all models and estimations.

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).

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).

Turning to the income-process parameters we consider two type of income processes. Table 1 displays the values for the parameters considered in each case.

In the RIP case, during working ages, we follow Kaplan and Violante (2010) and specify that income has a life cycle component, a persistent component, and an i.i.d component:

$$\log(y^i_{n,t}) = l(n) + z^i_{n,t} + \varepsilon^i_{n,t},$$

where $l(n)$ denotes the life cycle component, $\varepsilon^i_{n,t}$ is a transitory component, and $z^i_{n,t}$ is a persistent component that follows:

$$z^i_{n,t} = z^i_{n,t-1} + e^i_{n,t}.$$

We assume $\varepsilon^i_{n,t}$ and $e^i_{n,t}$ are normally distributed with variances σ^2_ε and σ^2_e, respectively.

In the HIP specification, during working ages, income has a life cycle component that is common to all households, a life cycle component that is idiosyncratic, a persistent component, and an i.i.d component:

$$\log(y^i_t) = l(n) + \alpha^i + \beta^i n + z^i_t + \varepsilon^i_t,$$

where $l(n)$ denotes the life cycle component common to all households of age n, $\alpha^i + \beta^i n$ is the life cycle component that is household-specific, z^i_t is a permanent component, and ε^i_t is a transitory component. As in Guvenen (2009), we assume the random vector (α^i, β^i) is distributed across households with zero mean, variances of σ^2_α and σ^2_β, and correlation of $corr_{\alpha\beta}$. Lastly, we assume the permanent component z^i_t follows an AR(1) process:

$$z^i_t = \rho z^i_{t-1} + e^i_t.$$

We assume ε^i_t, and e^i_t are normally distributed with variances σ^2_ε, and σ^2_e, respectively.

Importantly, we assume that borrowers and lenders are able to observe all components of income, including the household's (α^i, β^i). Once in retirement, the household receives a percentage of the last realization of the permanent component of its working-age income.

4.2 Estimation

Having assigned values to all first-stage parameters, we are in position to tractably estimate the remaining key parameters of interest. As is standard, we use a minimum distance estimator as in Chamberlain (1982, 1984), which minimizes a weighted squared sum of differences between model and data moments. The estimator solves the following problem:

$$\min_{\Theta} [\hat{g} - g(\Theta)]' W [\hat{g} - g(\Theta)], \tag{6}$$

where $g(\Theta)$ and $\hat{g}$ are $(J \times 1)$ vectors of model-based and data-based moments, respectively; and Θ is an $N \times 1$ vector of structural parameters to be estimated. W is a $J \times J$ weighting matrix, which we assume to be the identity matrix following Altonji and Segal (1996).

For all models we use the following sets of moments for identification:

1. Incidence of FD between ages 25-55 (30 moments).

Table 1: Parameters Determined Externally

Parameter	Value	Source
σ, Coefficient of relative risk aversion	2.0	Standard
r, Risk-free interest rate	3.0%	Standard
W, Retirement age	65	Standard
η, Roll-over interest rate on delinquent debt	20%	Livshits et al. (2007)
f, Bankruptcy filing cost (as a share of average income)	0.028	"
A_0, Replacement ratio	0.71	Hatchondo et al. (2015)
A_1, Replacement ratio	-0.045	"
A_2, Replacement ratio	0.14	"
ρ_n, mortality rate	–	Kaplan and Violante (2010)
RIP		
σ^2_ε, Variance of permanent shocks	0.05	Kaplan and Violante (2010)
σ^2_e, Variance of transitory shocks	0.01	"
HIP		
ρ, Autocorrelation of persistent shocks	0.821	Guvenen (2009)
σ^2_ε, Variance of persistent shocks	0.047	"
σ^2_e, Variance of transitory shocks	0.029	"
σ^2_α, Variance of intercept of life cycle income profile	0.022	"
σ^2_β, Variance of slope of life cycle income profile	0.00038	"
$corr_{\alpha\beta}$, Correlation between intercept and slope components	-0.23	"

2. Average persistence of FD at leads of 1-10 years (10 moments).

3. Concentration of FD at 1st-100th percentiles (100 moments).

4. Incidence of bankruptcy (BK) between ages 25-55 (30 moments).

5. Wealth-to-earnings ratio between ages 25-55 (30 moments).

These moments place strong constraints on what a successful model must replicate. First, a successful model must replicate salient facts on financial distress (e.g., its incidence, persistence, and concentration). Second, a successful model must account for the relative importance of informal delinquency versus formal bankruptcy when generating these aforementioned facts on financial distress. Lastly, all of this must be accomplished in the context where overall wealth accumulation patterns over the life cycle resemble those observed in the data.[17] Note that including patterns of wealth accumulation over the life cycle together with those of financial distress implies that we are asking the estimated model to generate sufficient cross-section variation in wealth.

[17]Our wealth-to-earnings moments are computed using data from the Survey of Consumer Finances between 1998-2016.

Notice that effectively we have 5 types of moments. Because each of these 5 set of moments differs in quantity (e.g., 30 moments summarizing the incidence of FD moments vs. 10 moments describing the persistence of FD moments) and in magnitude (e.g., the incidence of bankruptcy vs. wealth-to-earnings ratio), we make two adjustments. First, we "collapse" the dimensionality of the moments by weighting each moment by the inverse of the number of moments of the same type. For example, each incidence of FD moment is weighted by 1/30, whereas each persistence moment is weighted by 1/10. This first adjustment can be thought of as changing the weighting matrix to assign less weight to moments that belong to a group with many moments of the same type.[18] Second, we seek to minimize percentage deviations between data and model moments. In other words, $\hat{g} - g(\Theta)$ becomes $(\hat{g} - g(\Theta))/(0.5\hat{g} + 0.5g(\Theta))$. This second adjustment is also equivalent to changing the weighting matrix: we are effectively assigning less weight to a moment if it belongs to a type with a higher average level.

5 Results

5.1 Benchmark Model

This section discusses the estimation results for our benchmark model. First, Table 2 shows the fit of these models of two statistics for each type of target. Given that only three parameters were estimated, the model's fit is surprisingly good. For instance, rows (1) to (4) show that the model generates larger incidence of FD and bankruptcy for young individuals than for old individuals, as we observe in the data. The model also implies a significant rise in the average wealth-to-earnings ratio over the life cycle, reaching value for old individuals that are quite close to the data. Finally, and very importantly, these models also generate significant persistence and concentration of financial distress. In the next subsection, we will compare these predictions with the restricted models to determine what are the crucial features of the benchmark models. The key take-away from this table is that our benchmark models can simultaneously account for the incidence, persistence, and concentration of financial distress while generating reasonable patterns of wealth accumulation over the life cycle. Comparing the two benchmark models shows that the RIP specification generates more financial distress for older individuals than the HIP specification, and also a more reasonable wealth-to-earnings ratio for older individuals. More broadly, though, the results suggest the differences implied by the two income processes are quantitatively small.

[18]This reduction of moments is also used below to increase the power of the Sargan test, which follows Bowsher (2002), Roodman (2009), and Heathcote et al. (2014).

Table 2: Fit of Key Moments

	Data	DQBK β-het RIP	DQBK β-het HIP
(1) FD rate, age 25-34 (%)	15.0	19.8	21.7
(2) FD rate, age 35-44 (%)	13.1	11.3	10.6
(3) FD rate, age 54-55 (%)	10.0	8.36	7.05
(4) BK rate, age 25-34 (%)	0.87	1.51	1.90
(5) BK rate, age 35-44 (%)	1.00	0.76	1.00
(6) BK rate, age 45-55 (%)	0.78	0.60	0.67
(7) Wealth-to-earnings, age 25-34	1.12	0.57	0.63
(8) Wealth-to-earnings, age 35-44	2.04	1.84	1.98
(9) Wealth-to-earnings, age 45-54	3.29	3.56	3.96
(10) Average Pr(FD+3|FD)	0.42	0.38	0.40
(11) Average Pr(FD+5|FD)	0.28	0.29	0.29
(12) Average Pr(FD+8|FD)	0.17	0.22	0.21
(13) 70th percentile of FD	0.02	0.03	0.03
(14) 80th percentile of FD	0.12	0.17	0.16
(15) 90th percentile of FD	0.39	0.45	0.44

Columns (1) and (2) of Table 3 present the estimates values of the parameters. Looking at these two columns suggests that regardless of the income process used, the estimated parameter values are very similar. In other words, the data do not support the view that income-processes are easily discerned by the facts of financial distress. Put another way, financial distress is not driven, in any clear manner, by the structure of household income risk.

By contrast, we see that in both cases we can reject the null hypothesis that the low discount factor β_L equals the high discount factor β_H. This is of course an important point of our analysis: Purely homogeneous preferences do not appear to be selected by the data in the estimation.

5.2 Why Informal Default, and Why Discount Factor Heterogeneity?

Our preferred interpretation of the data is that it arises from a setting in which borrowers retain access to both formal (bankruptcy) and informal (delinquency) default, and in which borrowers vary in their subjective discount factors. To persuade the reader that this is the warranted inference, we now describe the performance of estimated alternatives and their implications. Specifically, in this section, we consider several alternative models to understand which features of our benchmark environment are key to delivering the facts.

First, to what extent do the data on FD allow us to distinguish between the two most prominent classes of income risk? To proceed, we first suppress discount-factor heterogeneity altogether, in

Table 3: Parameter values for estimated models

Parameter	DQBK β-het RIP (1)	DQBK β-het HIP (2)	DQBK No-het RIP (3)	DQBK No-het HIP (4)	No-DQ No-het RIP (5)	No-DQ No-het HIP (6)
Earnings threshold in DQ τ	7.464	6.600	4.363	15.171	–	–
	(1.747)	(1.612)	(0.650)	(50.00)		
Earnings threshold in BK τ	–	–	–	–	18.859	10.173
					(36.362)	(19.376)
Low discount factor β_L	0.807	0.784	0.846	0.930	0.932	0.936
	(0.065)	(0.058)	(0.040)	(0.017)	(0.015)	(0.018)
High discount factor β_H	$1.000^\dagger$	$1.000^\dagger$	–	–	–	–
Discharge shock to DQ debt γ	0.106	0.069	0.139	0.220	–	–
	(0.099)	(0.095)	(0.127)	(0.115)		
Share of pop. of type L	$0.500^\dagger$	$0.500^\dagger$	$1.000^\dagger$	$1.000^\dagger$	$1.000^\dagger$	$1.000^\dagger$

Notes: Asymptotic standard errors are in parenthesis. † denotes parameter fixed by assumption.

order to have a more direct assessment of the contribution of income dynamics to outcomes. to Table 4, Columns (3) and (4) present the estimated fit of these models with discount factor heterogeneity is removed (DQBK and no-het) while still allowing for either the RIP or HIP income processes.

The results suggest that once discount factor heterogeneity is suppressed, HIP income processes are favored over RIP processes within the DQBK model. Indeed, the χ^2 value of the DQBK model with the HIP process is 35 percent smaller than the DQBK model with the RIP process. As will be shown in the subsequent figures, in this case, allowing for ex-ante heterogeneity in the income process allows the model to replicate some of the incidence, persistence, and concentration of financial distress, while still generating a wealth-to-earnings profile that somewhat matches the data. Importantly, though, the DQBK HIP model in Column (4) produces a χ^2 value that is nearly five times larger than the equivalent model in Column (2) with discount factor heterogeneity suggesting the latter is a better description of the data than the former. This is an essential aspect of the estimation that leads us to conclude that discount-factor heterogeneity may well be a feature of the underlying environment.

Lastly, while the P-values in Columns (3) and (4) are considerably smaller than their counterparts in Columns (1) and (2), we cannot reject the null hypothesis that the over-identifying restrictions are satisfied by these models for commonly used levels of significance.[19] Indeed, the P-values suggest one can reject the aforementioned null hypothesis at the 15% level in the case of the No-het RIP model, and at the 29% level in the case of the No-het HIP model.

Turning next to the question of the extent to which informal default—by itself, i.e., again absent the presence of discount-factor heterogeneity—is a relevant option, Columns (5) and (6) in Table 4 present estimation results for models without discount factor heterogeneity and also without the delinquency margin (i.e., No-DQ and No-het). The results in those columns highlight the importance of modeling delinquency separate from formal bankruptcy. Indeed, the χ^2 value of the no-DQ RIP model (Column (5)) is over ten times larger than the equivalent figure from the benchmark RIP model (Column (1)). Similarly, though not as extreme, the χ^2 value of the no-DQ HIP model (Column (6)) is nearly six times larger than the equivalent statistic from the benchmark HIP model (Column (2)). Importantly, the P-value in Column (5) suggests one can reject the null hypothesis (that the over-identifying restrictions are satisfied) at the 10 percent

[19]This result is, in part, a function of our chosen weighting matrix. Estimating the models using a weighting matrix that penalizes more heavily deviations in the wealth-to-earnings moments would imply larger χ^2 values for the No-het models, with little change in the χ^2 values of the β-het models. As a result, it would be easier to reject the null under those circumstances.

level. Taken as a whole, these estimation results clarify the importance of allowing for both forms of income dynamics and both types of default.

In particular, we see that in the first two rows in Table 4, all models generate decreasing patterns for the incidence of financial distress over the life cycle, much like in the data. Once discount factor heterogeneity is dropped, however, both RIP and HIP models generate considerably less financial distress for older individuals than in the data. This failure becomes even more pronounced once the informal delinquency margin is dropped. In this case, both RIP and HIP models largely under-predict the incidence of financial distress of older individuals.

Table 4: Fit of Key Moments

	Data	DQBK β-het		DQBK No-het		No-DQ No-het	
		RIP	HIP	RIP	HIP	RIP	HIP
		(1)	(2)	(3)	(4)	(5)	(6)
(1) FD rate, age 25-34 (%)	15.0	19.8	21.7	19.4	23.6	15.4	13.9
(2) FD rate, age 35-44 (%)	13.1	11.3	10.6	7.82	8.30	4.79	5.28
(3) FD rate, age 54-55 (%)	10.0	8.36	7.05	5.59	3.15	0.71	2.31
(4) BK rate, age 25-34 (%)	0.87	1.51	1.90	1.61	1.29	-	-
(5) BK rate, age 35-44 (%)	1.00	0.76	1.00	0.54	0.53	-	-
(6) BK rate, age 45-55 (%)	0.78	0.60	0.67	0.42	0.28	-	-
(7) Wealth-to-earnings, age 25-34	1.12	0.57	0.63	0.19	0.40	0.31	0.45
(8) Wealth-to-earnings, age 35-44	2.04	1.84	1.98	0.09	0.81	0.33	0.93
(9) Wealth-to-earnings, age 45-54	3.29	3.56	3.96	0.12	1.80	1.06	2.08
(10) Average Pr(FD+3|FD)	0.42	0.38	0.40	0.30	0.22	0.05	0.08
(11) Average Pr(FD+5|FD)	0.28	0.29	0.29	0.20	0.14	0.04	0.08
(12) Average Pr(FD+8|FD)	0.17	0.22	0.21	0.14	0.08	0.03	0.05
(13) 70th percentile of FD	0.02	0.03	0.03	0.05	0.05	0.08	0.07
(14) 80th percentile of FD	0.12	0.17	0.16	0.17	0.19	0.25	0.22
(15) 90th percentile of FD	0.39	0.45	0.44	0.43	0.46	0.52	0.50
χ^2	-	0.52	0.53	3.90	2.53	5.49	3.10
P-value	-	0.77	0.77	0.14	0.28	0.06	0.22

Rows (3) and (4) of Table 4 show that all models that distinguish between delinquency and bankruptcy generate reasonable patterns for the incidence of bankruptcy over the life cycle. Because the initial wealth distribution and the filing cost are externally calibrated, all models in general imply higher filing rates than in the data for younger individuals. However, for older individuals the benchmark RIP and HIP models match the empirical target well.

More strikingly, Rows (5) and (6) of Table 4 show that only models with discount factor heterogeneity generate reasonable wealth accumulation patterns over the life cycle. Both the benchmark

RIP and HIP models imply sensible wealth-to-earnings ratios over the entire life-cycle. The ratios are low early in life, arising from life-cycle smoothing efforts, and high later in life, arising from planning for retirement. By contrast, once we leave the benchmark, wealth accumulation is wildly below that seen in the data—all models without discount factor heterogeneity imply grossly counterfactual wealth accumulation patterns. Additionally, we see that models with the RIP income process do worse than models with the HIP income process.

Another illuminating result is the stark contrast between models with and without discount factor heterogeneity in replicating the average persistence of financial distress. Rows (7) and (8) in Table 4 show that conditional on being distressed today, the average probability of being in distress 3 and 8 years out is very similar between the benchmark RIP and HIP models and the data. In contrast, models without discount factor heterogeneity systematically underpredict the persistence of financial distress at both horizons. Notably, the HIP model does considerably worse along this dimension than the RIP model. Finally, both No-DQ models imply very little persistence at both horizons. Overall, models without discount factor heterogeneity not only miss on the wealth targets, they also miss on the persistence of financial distress.

Finally, the last three Rows of Table 4 show that all models imply similar concentration of financial distress, with the No-DQ models implying slightly less. Both the benchmark and No-het models imply very similar Lorenz curves of financial distress regardless of the income process used. The No-het models imply less concentration, but are still able to generate the fact that most individuals are never in distress over the 18-year window in which they are observed. Indeed, for both the β-het and No-het models, the 70th percentile of financial distress is essentially zero, much like in the data.

In one sense, the findings above are natural, and flow from the stark fact of high persistence in financial distress in the data, in conjunction with the life-cycle needs of consumers. In particular, it is perhaps not surprising that when creditors are aware of borrowers' income risk (as assumed in essentially all the relevant literature)—and hence credit terms tighten as borrowers' conditions worsen, and when default of either the formal or informal kind carries stark consequences, financial distress would not be routinely utilized other than early in life, and not utilized often unless a borrower were substantially impatient.

6 Conclusion

This paper establishes first that using recently available proprietary panel data, while many US consumers (35%) experience financial distress as defined by severe (120 days past due) delinquency, at some point in the life cycle, most financial distress events are primarily accounted for 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 informal default and a simple form of heterogeneity in time preference, but not—within this fairly broad model class—otherwise. Specifically, the data are strongly consistent with the presence of a subset of effectively impatient consumers. We stress that the heterogeneity in effective discount factors that our estimation reveals is just that: Effective. Household behavior may well be rendered so potentially by a host of additional factors not modeled here. This implies that 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 via the (unwarranted) interpretation of our results as solely representing literal differences in time-preference.

References

Altonji, J. G. and L. M. Segal, "Small-Sample Bias In GMM Estimation Of Covariance Structures," *Journal of Business and Economic Statistics* 14 (July 1996), 353–366.

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, "Labor Market Upheaval, Default Regulations, and Consumer Debt," *Review of Economic Dynamics* 18 (January 2015), 32–52.

———, "Bankruptcy and delinquency in a model of unsecured debt," Technical Report, Federal Reserve Bank of St. Louis, 2017, forthcoming International Economic Review.

Athreya, K., X. Tam and E. Young, "Unsecured credit markets are not insurance markets," *Journal of Monetary Economics* 56 (2009), 83–103.

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.

Bowsher, C., "On Testing Overidentifying Restrictions in Dynamic Panel Data Models ," *Economic Letters* 77 (2002), 221–220.

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.

Chamberlain, G., "Multivariate Regression Models for Panel Data ," *Journal of Econometrics* 18 (1982), 5–46.

———, "Panel Data," in Z. Griliches and M. D. Intriligator, eds., *Handbook of Economometrics*volume 2 (North-Holland Press, 1984), 1247–1318.

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. Rios-Rull, "A Finite-Life Private-Information Theory of Unsecured Consumer Debt," *Journal of Economic Theory* 142 (2008), 149–177.

De Nardi, M., E. French and J. B. Jones, "Medicaid Insurance in Old Age," *American Economic Review* 106 (November 2016), 3480–3520.

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.

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 Econmics* 10 (Spring 1979), 33–54.

HEATHCOTE, J., K. STORESLETTEN AND G. L. VIOLANTE, " Consumption and Labor Supply with Partial Insurance: An Analytical Framework ," *American Economic Review* 104 (July 2014).

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.

Roodman, D. M., “A Note on the Theme of Too Many Instruments ,” *Oxford Bulletin of Economics and Statistics* 71 (2009), 135–158.

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 the 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 between the ages of 25 and 55 who enter the sample in 1999Q1.

Unconditional Fraction of Individuals in DQ. The unconditional fraction of individuals in delinquency (DQ), also called the 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 an individual is in delinquency at least one quarter at a particular age, $\mathbb{1}_{DQ_{i,j}} = 1$, the unconditional fraction of individuals in DQ is calculated as $\sum_{i=1}^{N_j} \mathbb{1}_{DQ_{i,j}}/N_j$.

Unconditional Fraction of Debt in DQ. Similarly, the 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, the unconditional fraction of debt in DQ is

$$\frac{\sum_{i=1}^{N_j} \frac{\text{DQ\_debt}_{i,j}}{\text{Total\_debt}_{i,j}}}{N_j}.$$

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

$$\frac{\sum_{i=1}^{N_j} \mathbb{1}_{DQ_{i,j}} \cdot \mathbb{1}_{DQ_{i,j+h}}}{\sum_{i=1}^{N_j} \mathbb{1}_{DQ_{i,j}}}.$$

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 the maximum age in the sample period is dropped. For example, if an individual i is in delinquency at age 40 in year 2014, $j^* = 43$, then 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 the Credit Limit. The unconditional probability of reaching the 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 at a particular age, $\mathbb{1}_{Credit_{i,j}} = 1$, then the unconditional probability of reaching the credit limit is calculated as

$$\frac{\sum_{i=1}^{N_j} \mathbb{1}_{Credit_{i,j}}}{N_j}.$$

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

$$\frac{\sum_{i=1}^{N_j} \mathbb{1}_{Credit_{i,j}} \cdot \mathbb{1}_{Credit_{i,j+h}}}{\sum_{i=1}^{N_j} \mathbb{1}_{Credit_{i,j}}}.$$

Average Life of DQ. The 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 the 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 spells 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 $DQnum_i$, the 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

$$\text{Delinquency Intensity} = \frac{\sum_{i=1}^{N_j} \frac{\text{DQ\_debt}_{i,j}}{\text{Total\_debt}_{i,j}}}{\sum_{i=1}^{N_j} \mathbb{1}_{DQ_{i,j}}}.$$

An alternative measure of delinquency intensity is calculated by taking the number of bankcards at least 120 DPD or Severe Derogatory. Let $\text{Num\_card}_{i,j}$ = crtr_attr17 + crtr_attr38, while $\text{Total\_card}_{i,j}$ = crtr_attr33 + crtr_attr34 + crtr_attr35 + crtr_attr36 + crtr_attr37 + 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 8 is computed by taking the 50th, 75th, and 90th percentiles 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 the 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 9 and 11 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 but so are the quantities. Thus, we see that across the US, financial distress patterns are very similar, and this is plausibly amenable to analysis within a model framework that abstracts from what might have seemed, a priori, as relevant differences across states.

Figure 9: The Persistence of Financial Distress Over the Life Cycle and Across States (debt)

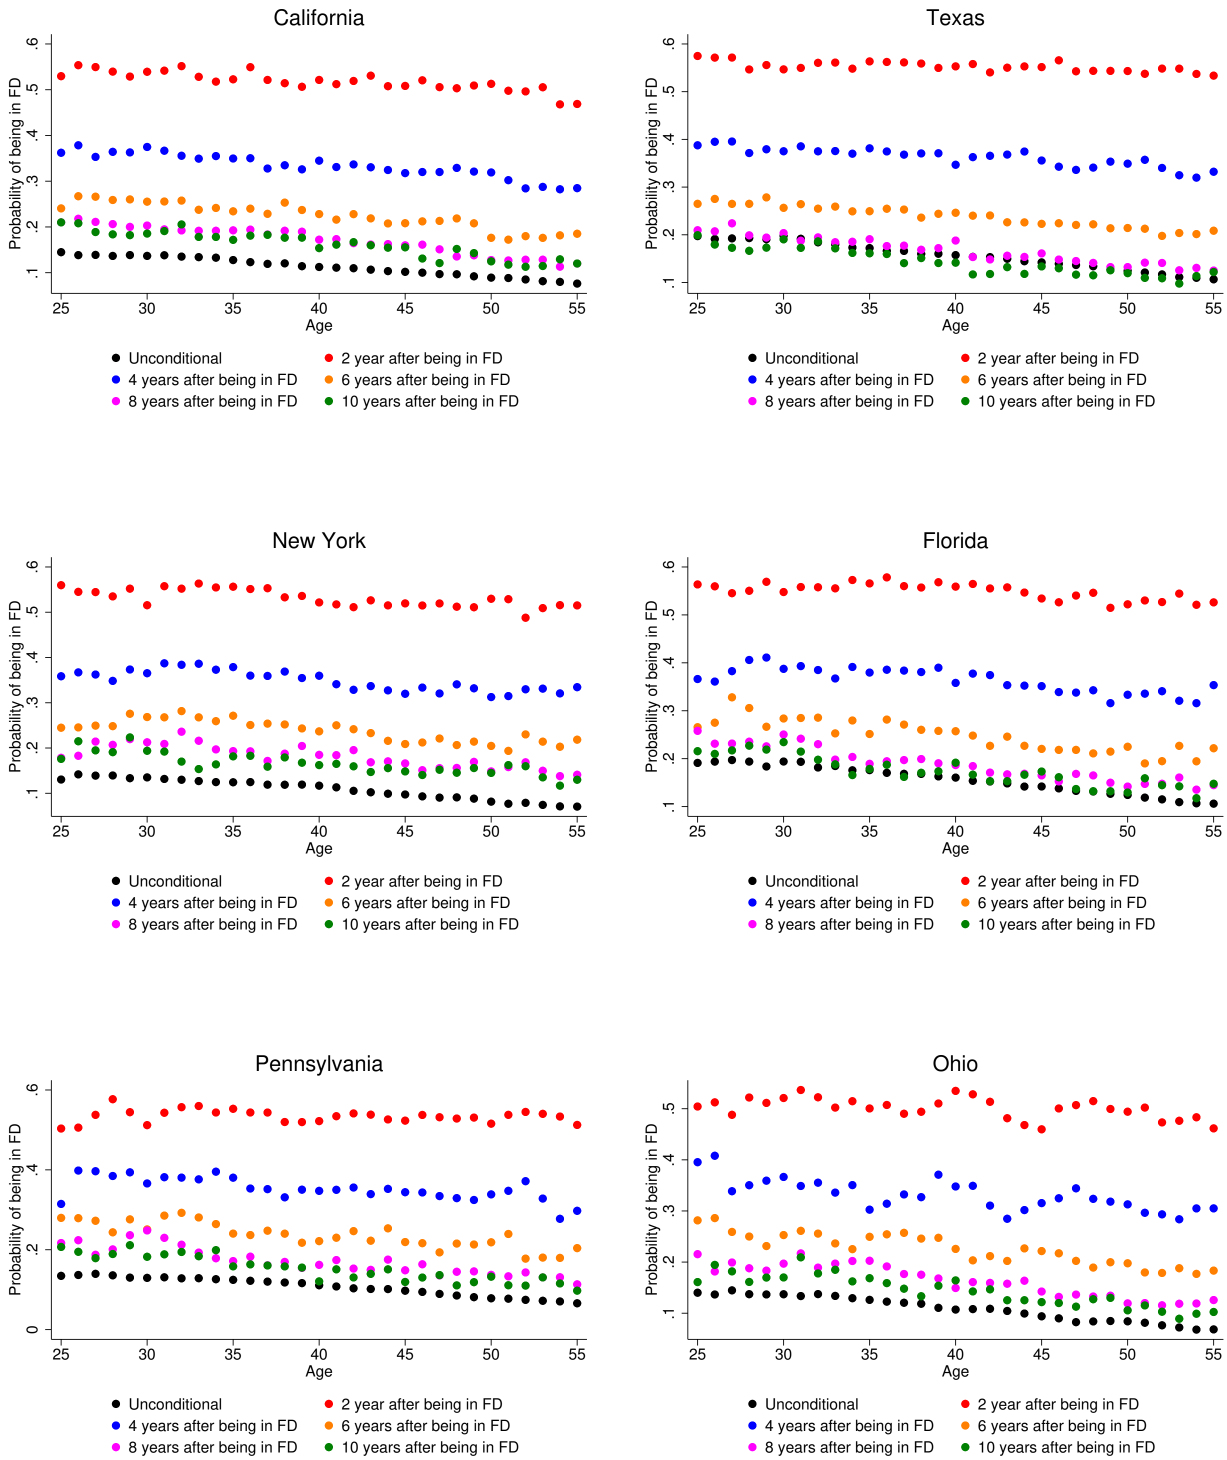

Source: Federal Reserve Bank of New York Consumer Credit Panel/Equifax

Figure 11: The Persistence of Financial Distress Over the Life Cycle and Across States (credit limit)

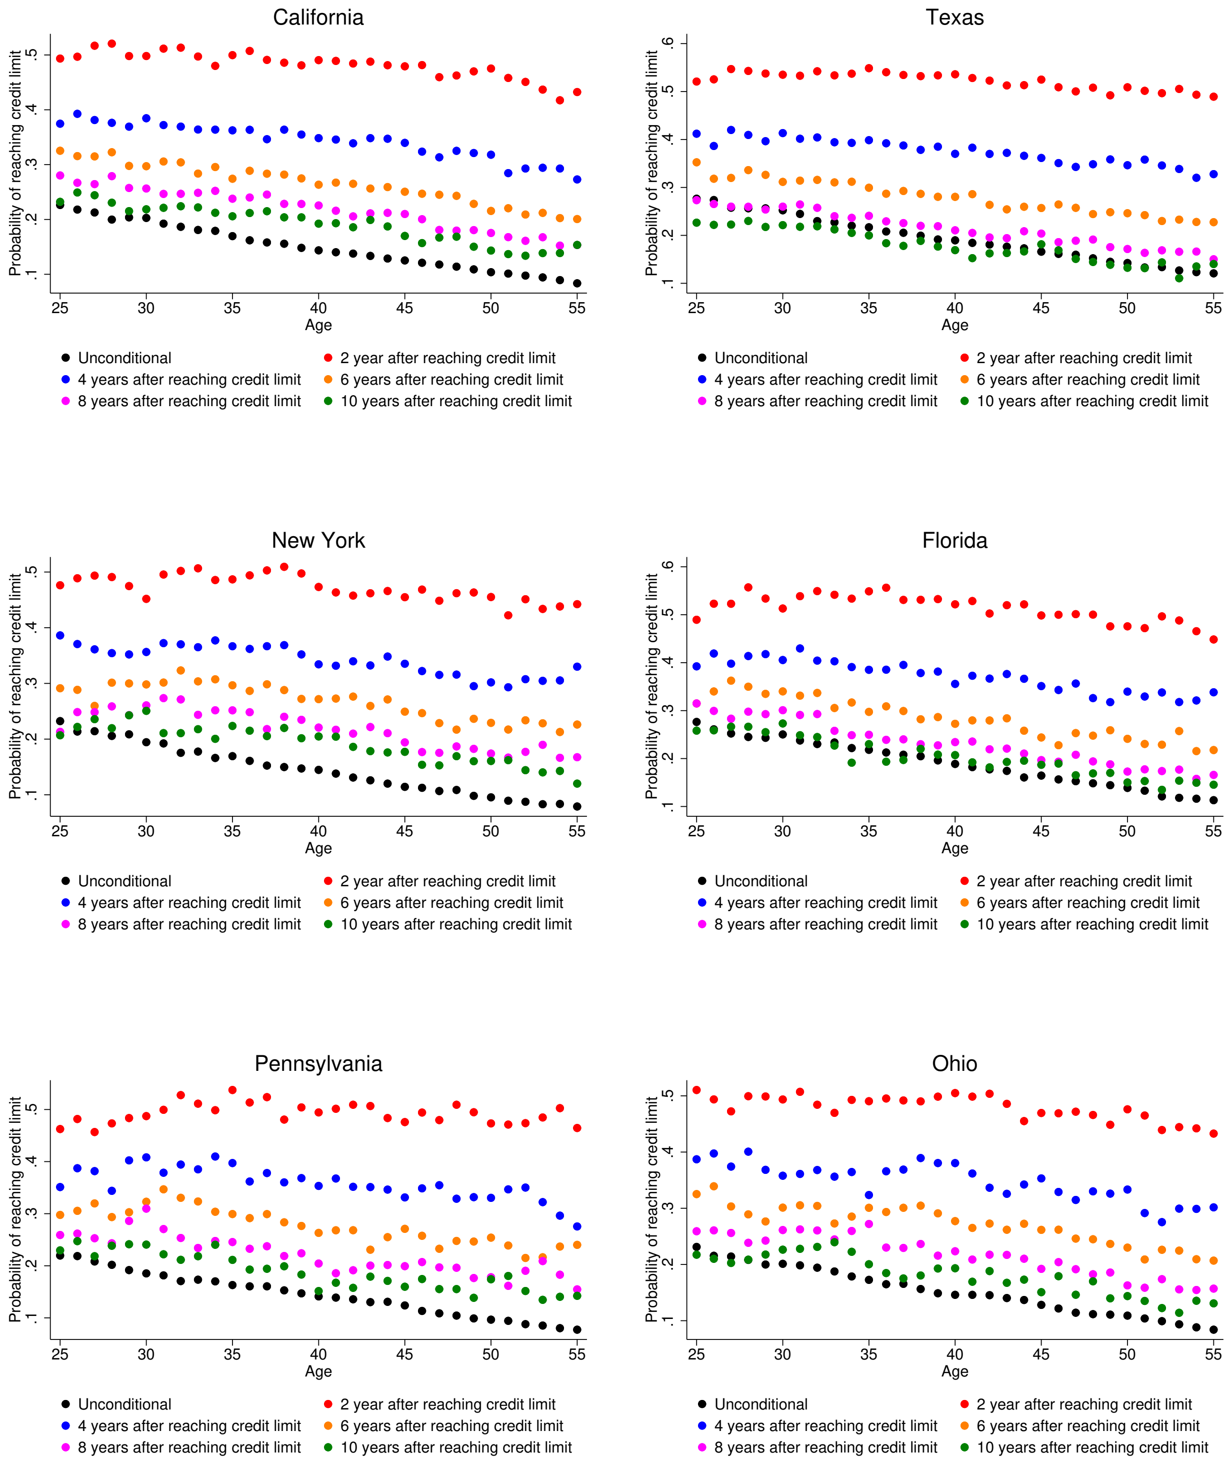

Source: Federal Reserve Bank of New York Consumer Credit Panel/Equifax

C A Model of Distress as Bankruptcy

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

We assume that in 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.[20] Additionally, in that period, consumption equals income—neither saving nor borrowing is allowed. In this environment, lifetime utility can be written as

$$G_{i,n}(z,\varepsilon,a) = \max\{\underbrace{V_{i,n}(z,\varepsilon,a)}_{\text{Pay}}, \underbrace{B_{i,n}(z,\varepsilon)}_{\text{Bankruptcy}}\} \tag{7}$$

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

$$B_{i,n}(z,\varepsilon) = u(\min\{y_{i,n}(z,\varepsilon),\tau\}) + \varrho_n \beta \mathbb{E}\left[G_{i,n+1}(z',\varepsilon',0)|z\right]. \tag{8}$$

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

[20] Chatterjee 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.

utility

$$V_{i,n}(z,\varepsilon,a) = \max_{\{a',c\}} u(c) + \varrho_n \beta \mathbb{E}\left[G_{i,n+1}(z',\varepsilon',a')|z\right],$$

$$\text{subject to} \tag{9}$$

$$c + a' q_{i,n}(z,a') = a + y_{i,n}(z,\varepsilon),$$

$$c \geq 0.$$

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.

$$q_{i,n}(z,a') = \frac{1}{1+r} \varrho_n \mathbb{E}\left[R_{i,n}(z',\varepsilon',a')|z\right]. \tag{10}$$

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 (7) to (9) are satisfied and prices satisfy the zero-profit condition (10).