

The Persistence of Financial Distress

Kartik Athreya

Federal Reserve Bank of Richmond

José Mustre-del-Río

Federal Reserve Bank of Kansas City

Juan M. Sánchez

Federal Reserve Bank of St. Louis

Using proprietary panel data, we show that many U.S. consumers experience financial distress (35% when distress is defined by having debt in severe delinquency, e.g.) at some point in their lives. However, most distress events are concentrated on a much smaller proportion of consumers in persistent trouble: 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 unsecured debt with informal default that accommodates a simple form of heterogeneity in time preference. (*JEL* D60, E21, E44)

Received November 10, 2017; editorial decision November 12, 2018 by Editor Stijn Van Nieuwerburgh. .

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 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 we tackle the second by estimating a battery of 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 is 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, that is, “severely delinquent.” Because severe delinquency is, in practice, an expensive way to repeatedly roll over debt, this definition plausibly captures financial distress. Put another way, because

We thank the Editor, Stijn Van Nieuwerburgh, and two anonymous referees for useful suggestions and guidance. 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. Send correspondence to Juan M. Sánchez, Federal Reserve Bank of St. Louis, PO Box 442, St. Louis, MO 63166. E-mail: juan.m.sanchez78@gmail.com.

© The Author(s) 2019. Published by Oxford University Press on behalf of The Society for Financial Studies. All rights reserved. For permissions, please e-mail: journals.permissions@oup.com.
doi:10.1093/rfs/hhz009

Advance Access publication February 1, 2019

severe 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. While lack of access to additional credit and the costs ensuing from failing to repay debt as promised are arguably the essence of financial distress, severe delinquency is not the only definition of FD one could use. Another measure of FD is one that tracks the extent to which consumers have depleted the credit lines available to them, and we will report on FD measured in this way as well.

As we will show, both measures we consider lead to a single, more general, conclusion: while many U.S. consumers (35% using 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, under the baseline definition, the incidence of FD is nearly double its unconditional rate 6 years after the initial distress event, and just 10% of borrowers account for half of all distress events. 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 facing households over a 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 not lasting very 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, experience distress. The latter is what we find in the data. Our empirical findings clearly indicate that the risk of financial distress is one that is, in a sense, resolved early in life: Most borrowers know that they will face few problems with timely debt repayment in the years ahead, and a much smaller few know that they will face a future of repeated instances of distress.

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 proceed with a formal structural estimation of a battery of alternative models of consumer debt and default and describe their implications. This analysis reveals 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 consumer debt with informal default that accommodates a simple form of heterogeneity in time preference.

By allowing for informal default, we capture an empirically relevant pathway for (non) repayment, as reflected by the substantial delinquency rates observed in U.S. data. By contrast, formal default (in its dominant “Chapter

¹ In addition to these facts, and as with the facts of the incidence described above, the persistence of financial distress is very similar across all 50 U.S. states. The interested reader is referred to the appendix.

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, MacGee, and Tertilt (2007), which was the first paper to allow for delinquency, or rollover at a “penalty” rate of interest, as an option in the period following a bankruptcy in case of an “expense” shocks.² As our results will show, informal default appears important in generating the observed persistence of financial distress at short horizons (1–3 years after the initial distress event).

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. Our allowance for preference heterogeneity is motivated by the prominent role it plays in a variety of models that are concerned with understanding the wealth distribution—often the extreme right tail and the heavy concentration one observes therein. Salient references here include Cagetti (2003), Krusell and Smith (1998, 1999), De Nardi and Fella (2017), and Favilukis, Ludvigson, and Nieuwerburgh (2017). As it will be clearer below (and in contrast to the role played by informal default in the persistence of near-term distress) discount factor heterogeneity also increases the persistence of financial distress for long horizons (8–10 years after the initial distress event).

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 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, MacGee, and Tertilt (2007); Chatterjee et al. (2007)). A consistent finding in this work (see also Athreya, Tam, and Young (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

² See Athreya et al. (2015) for a richer model of informal default, where, like 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.

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—such as any 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 by potentially the entire host of additional factors not modeled here. Future work that allows for more detail on household-level shocks, intrahousehold 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.

Financial distress or household financial “fragility” has received significant attention in recent work and has been the topic of interest with the general public (Picchi, 2015). Interest in the ability of the household to shield itself from susceptibility to shocks through the use of financial markets is, of course, longstanding. However, recent work has been aided by the arrival of more detailed data on household balance sheets (Lusardi, Schneider, and Tufano (2011), Lusardi (2011); Jappelli, Pagano, and Di Maggio (2013); Ampudia, van Vlokhoven, and Dawid (2016); Brunetti, Elena, and Costanza (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 questioning of how much borrowing the household could feasibly engage in, within a short-term period, for example, 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 United States as well as in the European Union 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 U.S. households that 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),

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 that 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.

Zeldes (1989), and others. Gross and Souleles (2002) use exogenous variation in credit line extensions to gauge the fraction that 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, Pistaferri, and Preston (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, MacGee, and Tertilt (2007) and Athreya (2008). For example, when distress is measured by severe delinquency (i.e., having a debt 120 days or more past due), models without delinquency or discount factor heterogeneity generate almost no persistence of financial distress at short horizons and very little at long horizons. Our quantitative results show that delinquency is useful in correcting the former, while discount factor heterogeneity helps with the latter.

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.⁵ 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

⁵ 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). See Frederick, Loewenstein, and O’Donoghue (2002) for a survey.

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. In turn, this 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 U.S. households substantially vary 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. Third, while not about consumer credit use, Parker (2017) finds that U.S. 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).⁶ 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.

1. Financial Distress in the United States

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 U.S. account holders.

We focus on individuals with credit histories between 1999Q1 and 2016Q4 and who enter our observation frame between the ages of 25 and 55. Additionally, we restrict our attention to individuals who have at least 10 years, worth of credit data. Because our model will focus on default and delinquency behavior prior to retirement, we further restrict our measurements to individuals through the age of 65.⁷ Combining these restrictions implies, for example, that

⁶ Relatedly, Mustre-del Río (2015) finds that persistent employment differences across males in the United States cannot be explained by differences in wealth or wages and hence are indicative of persistent differences in the disutility of work.

⁷ All figures plot data through age 55, because we measure default up to 10 years in the future.

65-year-olds entered our sample as early as 1999 (at the age of 48) and as late as 2007 (at the age of 56).

While our analysis focuses on a specific group of individuals 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 1999Q1–2016Q4 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. As stated earlier, we define an individual as being 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. The appendix offers additional details about our data.

In this section, for the figures that follow, we ask the reader to focus on the left panels, which represent the data. We will return to compare the former with the right panels—which represent the predictions of our preferred model—after we formally describe the model(s) and their estimation. To start, consider the following question: How broadly shared an experience is financial distress? The left panel of Figure 1 takes a life-cycle perspective. The solid dots show 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 dots in the figure begin near 14% among the young and fall below 10% later in life.

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 the left panel of Figure 1 we see very clearly just how persistent the state of financial distress is for U.S. consumers. Conditional on being in distress today, the likelihood of being distressed in 6 years⁸ (the solid squares) is nearly double that of the unconditional rate (the solid dots) over the entire life cycle. As we will further show below, this particular feature will elude the standard model of defaultable consumer debt and will instead suggest the importance of delinquency and heterogeneity in individual time preference.

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 as 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 it 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

⁸ 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 6 years.

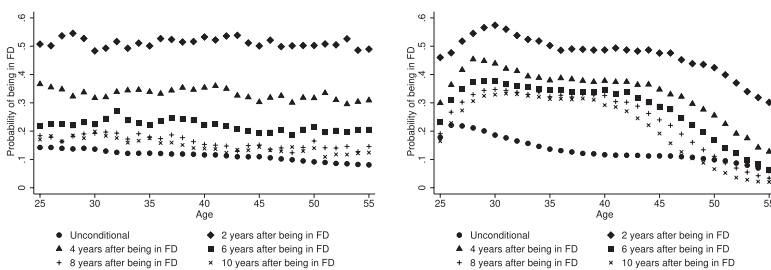


Figure 1
The persistence of FD over the life cycle (debt): Data (left) and model (right)
 Source: See the appendix.

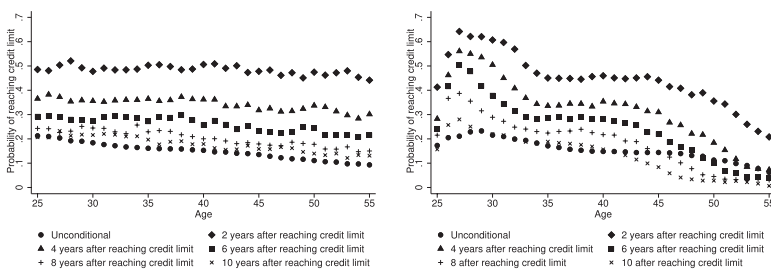


Figure 2
The persistence of FD over the life cycle (credit limit): Data (left) and model (right)
 Source: See the appendix.

cited at the outset: that of a large subpopulation being unable to raise funds in an emergency. The left panel of Figure 2 (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.⁹

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 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).

The left panel of Figure 3 provides further evidence on distress, this time measured by the proportion of time in the sample that a consumer spends in distress. To improve the exposition, this figure excludes those who experience

⁹ In our data set, on average, about 50% of individuals in delinquency have also depleted their credit.




Figure 3
Share of time in FD: Data (left) and model (right)
 Source: See the appendix.




Figure 4
Number of FD spells: Data (left) and model (right)
 Source: See the appendix.

no distress during their time in the sample. The left panel of Figure 3 shows that while distress is indeed fleeting for some (roughly 35%), for 65% of consumers, distress is a much more routine state of affairs. Among those who experience financial distress at least once, more than 40% of them are distressed for at least a quarter of the time for which we observe them.

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. The left panel of Figure 4 shows how, for those who have experienced financial distress at least once, the number of spells they experience is often substantial, with roughly one-tenth of the sample experiencing three or more spells.

Our measures so far all imply that financial distress is “concentrated” on the population. Formally, perhaps the most natural way to demonstrate this is via the Lorenz curves presented in the left panel of Figure 5. 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 (solid line) or having depleted available credit (dashed line).




Figure 5
The concentration of financial distress: Data (left) and model (right)
 Source: See the appendix.

These facts suggest that financial distress may well be an important phenomenon, especially from the point of view of a subset of individuals looking out, *ex ante*, over a life cycle.¹⁰ Discerning this importance, however, requires specifying and evaluating models, which we turn to in the next sections. 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.

2. 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, MacGee, and Tertilt (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 (1) the choice to informally default at any time and (2) a simple form of heterogeneity in time preference. Using standard techniques, we estimate the key parameters of our benchmark model and show these extensions are sufficient to account for the facts. While establishing the absolute necessity of these features is not easy, we will show that when a fairly comprehensive set of alternative models is estimated to match the same set of facts, these models fail to simultaneously generate the incidence, persistence, and concentration of financial distress—with a key sticking point being the ability to generate

¹⁰ The appendix establishes 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.

empirically reasonable patterns of FD at short and long horizons, as well as wealth accumulation over the life cycle.

2.1 A benchmark 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 q_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.¹¹ Moreover, in any period of delinquency, consumption equals income up to a threshold τ .¹² Second, 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, income garnishment does not happen 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 is potentially more patient and thus has a discount factor $\beta_H \geq \beta_L$. Denote an individual's discount type by j .

In this framework, lifetime utility is written as

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

where V , B , and D are lifetime utilities for households paying back their debt, filing for bankruptcy, and being delinquent on their debt, respectively. Note that these functions are indexed by a household's discount factor type j and age n . These functions take as arguments current household wealth/debt and the household's income state. The latter is summarized by a permanent component

¹¹ This representation captures the main ingredients in Athreya et al. (2018). They show evidence that penalty rates modeled this way are able to capture the key features of delinquency.

¹² The remaining income is lost, for instance, when dealing with debt collectors.

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

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

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

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

with $a' = (1 + \eta)a.$ (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

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

subject to (4)

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

$$c \geq 0,$$

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:

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

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:

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

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).

3. Calibration and Estimation

Our approach to model parameterization is standard: We assign parameter values in a two-step procedure.¹³ 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 assume that type H individuals have a discount factor of $\beta_H = 1$.¹⁴ This leaves us a total of four parameters ($\tau, \beta_L, p_L, \gamma$) to be estimated

¹³ See, for example, De Nardi, French, and Jones (2016).

¹⁴ Appendix Appendix D shows that our main results are unchanged when we relax this assumption as we cannot reject $\beta_H = 1.00$ in any estimation.

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 a battery of alternatives to our benchmark model. These alternative models fall into two groups. 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. The empirical performance of these No-het models compared to our benchmark models helps us assess the importance of discount factor heterogeneity within the context of models with a delinquency margin. Second, to assess the importance of the delinquency margin, we also drop the informal delinquency margin and consider models with only bankruptcy (referred to as No-DQ). Here too, we still consider both RIP and HIP income processes, as well as models with and without discount factor heterogeneity (these models are described in Appendix Appendix F). The empirical performance of these No-DQ models compared to our benchmark and No-het models helps us assess the importance of modeling informal delinquency separate from formal bankruptcy.

3.1 Assigning first-stage parameters

Across all models, a period in the model refers to 1 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 at 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 at 20% annually, following Livshits, MacGee, and Tertilt (2007). Bankruptcy filing costs are at 2.8% of average

income, or roughly \$1,000, again following Livshits, MacGee, and Tertilt (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: $\max\{A_0 + A_1 \exp(z_{i,w-1}), A_2\}$. To be consistent with U.S. 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 types of income processes. The bottom panels of Table 1 display the parameter values 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:

$$\log(y_{n,t}^i) = l(n) + z_{n,t}^i + \varepsilon_{n,t}^i,$$

where $l(n)$ denotes the life-cycle component, $\varepsilon_{n,t}^i$ is a transitory component, and $z_{n,t}^i$ is a persistent component that follows:

$$z_{n,t}^i = z_{n,t-1}^i + e_{n,t}^i.$$

We assume $\varepsilon_{n,t}^i$ and $e_{n,t}^i$ are normally distributed with variances σ_ε^2 and σ_e^2 , 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:

$$\log(y_t^i) = l(n) + A^i + B^i n + z_t^i + \varepsilon_t^i,$$

where $l(n)$ denotes the life-cycle component common to all households of age n , $A^i + B^i n$ is the life-cycle component that is household-specific, z_t^i is a permanent component, and ε_t^i is a transitory component. As in Guvenen (2009), we assume the random vector (A^i, B^i) is distributed across households with zero mean, variances of σ_A^2 and σ_B^2 , and correlation of $corr_{A,B}$. Lastly, we assume the permanent component z_t^i follows an AR(1) process:

$$z_t^i = \rho z_{t-1}^i + e_t^i.$$

We assume ε_t^i , and e_t^i are normally distributed with variances σ_ε^2 , and σ_e^2 , respectively.

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

Table 1
Externally determined parameters

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 , Rollover interest rate on delinquent debt	20%	Livshits, MacGee, and Tertilt (2007)
f , Bankruptcy filing cost (as a share of average income)	0.028	"
A_0 , Replacement ratio	0.71	Hatchondo, Martinez, and Sánchez (2015)
A_1 , Replacement ratio	-0.045	"
A_2 , Replacement ratio	0.14	"
ρ_n , mortality rate	-	Kaplan and Violante (2010)
RIP		
σ_ε^2 , Variance of permanent shocks	0.05	Kaplan and Violante (2010)
σ_e^2 , Variance of transitory shocks	0.01	"
HIP		
ρ , Autocorrelation of persistent shocks	0.821	Guvenen (2009)
σ_ε^2 , Variance of persistent shocks	0.047	"
σ_e^2 , Variance of transitory shocks	0.029	"
σ_A^2 , Variance of intercept of life-cycle income profile	0.022	"
σ_B^2 , Variance of slope of life-cycle income profile	0.00038	"
$corr_{A,B}$, Correlation between intercept and slope components	-0.23	"

3.2 Estimation

Having assigned values to all first-stage parameters, we are in a position to tractably estimate the remaining key parameters of interest. As is standard, we use a minimum distance estimator like in Chamberlain (1982, 1984). Their estimator minimizes a weighted squared sum of differences between the model and the 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 the ages 25 and 55 (30 moments);
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 the ages 25 and 55 (30 moments);
5. Wealth-to-earnings ratio between the ages 25 and 55 (30 moments);
6. Wealth Gini index (1 moment); and
7. Interquartile range of wealth distribution (as a fraction of median wealth) (1 moment).

These moments place strong constraints on what a successful model must replicate. First, a successful model must replicate salient facts on financial distress (especially 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.¹⁵ Note that including wealth inequality moments together with those of financial distress implies that we are asking the estimated model to generate sufficient cross-section variation in wealth.

Notice that effectively we have 7 sets of moments. Because each of these 7 sets of moments differ 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.¹⁶ Second, we seek to minimize percentage deviations between the data and the 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 set with a higher average level.

4. Results

4.1 Benchmark model

This section discusses the estimation results for our benchmark model. Columns 1 and 2 of Table 2 summarize the fit of our benchmark model (DQBK β -het) by income process used. Given that only four parameters were estimated, the model’s fit is surprisingly good. For instance, rows 1 to 6 show that the model generates higher incidences of FD and bankruptcy at younger compared to older age, as we observe in the data. The model also implies a significant rise in the average wealth-to-earnings ratio over the life cycle, and it generates wealth inequality that is quite close to the data. Finally, and very importantly, these models also generate significant persistence and concentration of financial distress.

¹⁵ Our wealth-to-earnings moments are computed using data from the Survey of Consumer Finances between 1998 and 2016.

¹⁶ This reduction of moments is also used below to increase the power of the Sargan test, which follows Bowsher (2002), Roodman (2009), and Heathcote, Storesletten, and Violante (2014).

Table 2
Fit of key moments

Data	DQBK β -het		DQBK No-het		No-DQ β -het		No-DQ No-het		
	RIP	HIP	RIP	HIP	RIP	HIP	RIP	HIP	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	
(1) FD rate, ages 25–34 (%)	13.42	21.07	21.30	22.18	25.16	14.28	13.66	18.07	14.84
(2) FD rate, ages 35–44 (%)	11.69	13.59	12.07	8.99	9.69	10.14	7.21	8.65	6.28
(3) FD rate, ages 45–55 (%)	9.35	10.87	9.29	5.91	4.06	9.67	5.94	4.20	3.04
(4) BK rate, ages 25–34 (%)	0.87	1.39	1.52	1.56	1.24	–	–	–	–
(5) BK rate, ages 35–44 (%)	1.00	0.74	0.64	0.62	0.55	–	–	–	–
(6) BK rate, ages 45–55 (%)	0.78	0.71	0.61	0.46	0.33	–	–	–	–
(7) Wealth-to-earnings, ages 25–34	1.12	0.66	0.80	0.20	0.37	0.70	0.76	0.27	0.40
(8) Wealth-to-earnings, ages 35–44	2.04	2.20	2.66	0.10	0.66	2.34	2.37	0.19	0.72
(9) Wealth-to-earnings, ages 45–55	3.29	4.27	5.33	0.15	1.46	4.53	4.72	0.42	1.59
(10) Wealth Gini	0.76	0.70	0.66	0.86	0.68	0.68	0.69	0.73	0.68
(11) Wealth ($P75 - P25$)/ $P50$	3.62	5.81	5.02	4.10	3.27	4.46	7.82	3.39	3.48
(12) Average Pr(FD+1 FD)	0.69	0.63	0.65	0.56	0.50	0.00	0.00	0.00	0.00
(13) Average Pr(FD+3 FD)	0.41	0.39	0.41	0.28	0.24	0.15	0.12	0.08	0.09
(14) Average Pr(FD+5 FD)	0.27	0.31	0.33	0.19	0.15	0.21	0.14	0.07	0.09
(15) Average Pr(FD+8 FD)	0.16	0.24	0.25	0.12	0.09	0.24	0.13	0.05	0.06
(16) Average Pr(FD+10 FD)	0.15	0.22	0.22	0.10	0.07	0.21	0.11	0.04	0.05
(17) 70th percentile of FD	0.08	0.06	0.04	0.08	0.08	0.10	0.10	0.25	0.13
(18) 80th percentile of FD	0.23	0.23	0.18	0.23	0.25	0.30	0.27	0.43	0.29
(19) 90th percentile of FD	0.52	0.51	0.47	0.49	0.51	0.59	0.55	0.66	0.56
χ^2		0.83	0.96	3.68	2.26	0.78	1.54	5.14	2.78
p -value		.84	.81	.45	.69	.94	.82	.40	.73

More specifically, the RIP specification (Column 1) generates slightly more financial distress for older individuals than does the HIP specification, and it also generates more reasonable wealth-to-earnings ratios for older individuals.

Columns 1 and 2 of Table 3 present the estimated parameter values. Looking at these two columns shows 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. That is, financial distress is not driven, in any clear manner, by either of the two standard representations (HIP, RIP) 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 . Of course, this is an important point of our analysis: purely homogeneous preferences do not appear to be selected by the data in the estimation.¹⁷

¹⁷ One may wonder at this point how much more heterogeneity in time preference is needed to fit FD, in addition to what is needed to fit wealth inequality. We answered that question reestimating our benchmark model eliminating moments about concentration and persistence of financial distress in Appendix Appendix E. We conclude from this exercise that while both wealth inequality and persistence/concentration FD “ask” for some degree of beta heterogeneity to reproduce the data, wealth inequality “asks” for even larger differences in discounting across consumers.

Table 3
Parameter values for estimated models

Parameter	DQBK β -het		DQBK No-het		No-DQ β -het		No-DQ No-het	
	RIP (1)	HIP (2)	RIP (3)	HIP (4)	RIP (5)	HIP (6)	RIP (7)	HIP (8)
Earnings threshold in DQ τ	29.837 (13.717)	23.451 (8.423)	4.965 (0.649)	24.387 (10.574)	—	—	—	—
Earnings threshold in BK τ	—	—	—	—	29.513 (11.609)	12.870 (3.219)	27.213 (3.374)	9.815 (1.196)
Low discount factor β_L	0.809 (0.049)	0.796 (0.049)	0.859 (0.032)	0.930 (0.017)	0.728 (0.059)	0.827 (0.053)	0.902 (0.008)	0.926 (0.012)
High discount factor β_H	1.000 [†]	1.000 [†]	—	—	1.000 [†]	1.000 [†]	—	—
Discharge shock to DQ debt γ	0.143	0.142	0.139	0.232	—	—	—	—
Share of pop. of type L	0.401 (0.044)	0.327 (0.054)	1.000 [†]	1.000 [†]	0.361 (0.013)	0.408 (0.070)	1.000 [†]	1.000 [†]

Asymptotic standard errors are reported in parentheses. [†] denotes a parameter fixed by assumption.

To look at the moments of FD more closely, return now to Figures 1-5 discussed at the outset. Recall that the left panel in each figure shows the data, while the right panel shows the corresponding measurement from our preferred model: DQBK β -het RIP.

First, Figure 1 shows that (our preferred specification) can generate a life-cycle profile of the incidence of financial distress (solid dots) that is quite close to the data. Figure 2 shows that this model can also recreate the life-cycle incidence of reaching one’s credit limit, which was an outcome not directly targeted by the estimation.¹⁸

The remaining markers in Figures 1 and 2 also show that our preferred model generates life-cycle profiles of the persistence of financial distress that are quite similar to the data. Since the model does not have a reason to be in distress after retirement, the fit of the data is more challenging for older individuals.¹⁹

Next, Figure 3 shows the distribution of the time in the sample that individuals spent in FD. Recall that for this figure we excluded those who had no FD during their time in the sample.²⁰ Although this distribution was not targeted at all in the estimation, the model generates a surprisingly good fit of the data. For instance, in the data, about 40% of the individuals who had at least one event of FD spent at least a quarter of the time in the sample in FD. In the model, the same share is slightly less than 50%.

¹⁸ In the model, we define the “credit limit” as the level of debt in which the implied borrowing rate is higher than 16%. That rate was chosen to match the average incidence (across all ages) of people reaching the credit limit in the data.

¹⁹ Although this issue is clearly beyond the scope of this paper, we believe that adding health shocks, which are more important for older individuals, may help reconcile the model with the data.

²⁰ In the model, we create a sample with the same characteristics used in the data to generate this histogram. In particular, the age distribution of model-based data matches that of the data. Additionally, conditional on age, the distribution of time (years) in the sample matches that of the data.

Also not targeted is the number of spells in FD, presented in Figure 4. The model has a reasonable fit as well. For instance, the share with one or two spells is almost 90% in the data and almost 80% in the model.

Finally, Figure 5 shows that the model generates a concentration of FD that is quite similar to the data. The main difference is that while in the data there is a higher concentration of FD defined using delinquency, the model implies there is slightly higher concentration for individuals that reached their credit limit.

In the next subsection, we will compare these predictions with the restricted models to determine what are the crucial features of our benchmark model. The key takeaway from this section is that our benchmark model can simultaneously account for the incidence, persistence, and concentration of financial distress while generating reasonable patterns of wealth accumulation over the life cycle.

4.2 Why informal default, and why discount factor heterogeneity?

Our preferred interpretation of the data is that they arise from a setting in which borrowers retain access to both formal (bankruptcy) and informal (delinquency) default and vary in their subjective discount factors (beta-heterogeneity). 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, consider a standard model of default with no beta-heterogeneity and no delinquency, as presented in Columns 7 and 8 of Table 2.²¹ In the standard model of default, with no beta-heterogeneity and no delinquency, after an episode of FD the individual gets a fresh start and exits FD with no debt (as in a Chapter 7 bankruptcy). In the years immediately following FD, although income is still relatively low, another episode of FD is unlikely because debt is very low as well. Eight or ten years after FD, this individual may still have relatively low income (at least in the RIP case, in which shocks are extremely persistent), but he/she is unlikely to have a significant amount of debt, because in this model most debt is accumulated for life-cycle reasons and this previously distressed individual is now older. This explains why in the standard model of FD presented in Columns 7 and 8 there is almost no persistence of FD.

Next, consider Columns 5 and 6 of Table 2, which show the results of adding beta-heterogeneity. Here, we note that beta-heterogeneity by itself does not increase near-term FD (i.e., 1–3 years immediately following the initial FD episode), but it does increase long-term FD (i.e., 8–10 years after the initial FD episode). This can be seen for both types of income processes, but the finding

²¹ In all the models without delinquency we interpret the option of discharging all debt as FD instead of bankruptcy. As can be seen in Columns 7 and 8 of Table 2, this model replicated rates of FD that are very large compared to the BK rates in the data. Appendix Appendix G shows that the results are robust to considering BK rates' moments in the estimation instead of FD rates.

is more salient for the RIP-income process specification (cf. Columns 5 and 7), where income shocks are more persistent. The main reason for this success at long horizons is that individuals with a low discount factor accumulate debt even when they are older.

Separately, consider Columns 3 and 4 of Table 2, which show the model with delinquency but without beta-heterogeneity. Here, the main finding is that adding delinquency by itself increases near-term FD, but not long-term FD. This can be seen by comparing Columns 3 and 4 with 7 and 8, respectively. This success occurs because delinquency allows for an additional margin by which agents can accumulate debt and remain in FD in the years immediately following the initial FD episode. However, in the long-term, life-cycle forces dominate and individuals do not have a significant amount of debt and hence are unlikely to enter into FD 8–10 years out.

With these decompositions in hand it is not surprising that models with both beta-heterogeneity and delinquency (Columns 1 and 2) can generate both near-term FD and long-term FD. The former is mostly due to the delinquency margin, whereas the latter is mostly due to discount factor heterogeneity.

5. Conclusion

This paper establishes first that using recently available proprietary panel data, while many U.S. 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—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. 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.

Appendix 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 focus on individuals with credit histories between 1999Q1 and 2016Q4 and who enter our observation frame between the ages of 25 and 55. Additionally, we restrict our attention to

individuals who have at least 10 years of credit data. Finally, we further restrict our measurements to individuals through the age of 65. Thus, these restrictions imply, for example, that 65-year-olds entered our sample as earlier as 1999 (at the age of 48) and as late as 2007 (at the age of 56). Similarly, 25-year-olds could have entered as late as 2007 and exit in 2016 at the age of 34, to meet the 10-year minimum requirement.

A1. 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; that is, $DQ_debt_{i,j} = crtr_attr111 + crtr_attr112$, 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 $DQ_debt_{i,j} > 0$. 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 $\sum_{i=1}^{N_j} \mathbb{1}_{DQ_{i,j}} / N_j$.

A2. 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; that is, $Total_debt_{i,j} = crtr_attr107 + crtr_attr108 + crtr_attr109 + crtr_attr110 + crtr_attr111 + crtr_attr112$. Then, the unconditional fraction of debt in DQ is

$$\frac{\sum_{i=1}^{N_j} \frac{DQ_debt_{i,j}}{Total_debt_{i,j}}}{N_j}$$

A3. 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.

A3. Unconditional Probability of Reaching the Credit Limit The unconditional probability of reaching the credit limit is calculated by finding the ratio of individuals reaching the 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; that is, $crtr_attr169 \geq crtr_attr180$. 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

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

A4. Conditional Probability of Reaching the Credit Limit

Similarly, the probability of reaching the credit limit, conditional on reaching the credit limit h years ago, is computed as

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

A5. Average Time in DQ

The average time in DQ for individual i is computed as the ratio of total number of years i is in DQ ($\mathbb{1}_{DQ_{i,j}} = 1$) to the total number of years in the sample for i . Recall that individuals can be in the sample for at least 10 years and at most 18. Additionally, note that an individual is considered in DQ in a given year if she is in DQ for at least one quarter of that year. Let $DQnum_i$ be the total number of years i is in DQ, and let T_i denote the total number of years in the sample period for i . Then

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

Note that Figure 3 excludes individuals who do not spend any time in DQ because the large proportion of the population that does not enter DQ distorts the scale of the histogram.

A6. Delinquency Spell Number

A delinquency spell begins when the individual is in DQ ($\mathbb{1}_{DQ_{i,j}} = 1$) in the current year but was not in DQ the preceding year. Similarly, a delinquency spell ends when the individual is not in DQ in a year but was in DQ the preceding year. If the first and last observation is in DQ, we take that year to be the start or end of the DQ, respectively. Note that an individual can have multiple delinquency spells throughout her life.

A7. Lorenz Curves

Lorenz curves are calculated using two measures: being in DQ and reaching the 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.

A8. Delinquency Intensity

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

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

An alternative measure of delinquency intensity is calculated by taking the number of bankcards at least 120 DPD or Severe Derogatory. Let $\text{Num_card}_{i,j} = \text{crtr_attr17} + \text{crtr_attr38}$, 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 $\mathbb{1}_{\text{Card}_{i,j}} = 1$ if $\text{Num_card}_{i,j} > 0$. Then it is computed as

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

A9. Delinquent Debt

Figure B2 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 January 2017 dollars using seasonally adjusted CPI from the U.S. Bureau of Labor Statistics.

Appendix B. Additional Figures

The measures presented in the main text are primarily “extensive margin” measures: They are based on measures of financial distress that are binary—whether someone has severely delinquent debt, or 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 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 B1 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 line), 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 85% at older ages). A similarly flat life-cycle profile is observed when intensity is measured as the average number of accounts severely delinquent (dashed line). Overall, this figure suggests that when individuals are categorized as in financial distress based on our extensive margin measure, it is because most of their debt and accounts are severely delinquent.




Figure B1
Intensity of FD




Figure B2
Distressed debt

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

What about the size distribution of distressed debts? Figure B2 summarizes the amount of delinquent debt by the 50th, 75th, and 90th percentiles by age.²² 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.

Collecting the extensive- and intensive-margin empirics, we can summarize our findings as follows: Financial distress among U.S. 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.

Appendix C. 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 C1 and C2 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 United States, 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.

Appendix D. Estimating β_H

In this section we relax the assumption that $\beta_H = 1.00$ and allow this parameter to be estimated as well. Table D1 shows that in all cases the resultant point estimates and standard errors for β_H are such that we cannot reject the hypothesis that the estimated value equals 1.00 at any reasonable degree of significance.

Next, the results in Table D2 reveal that the findings from the main text are also unchanged when we allow β_H to be estimated. Indeed, we still find that delinquency is needed to generate near-term persistence of financial distress: the No DQ models in Columns 3 and 4 generate very little

²² 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).

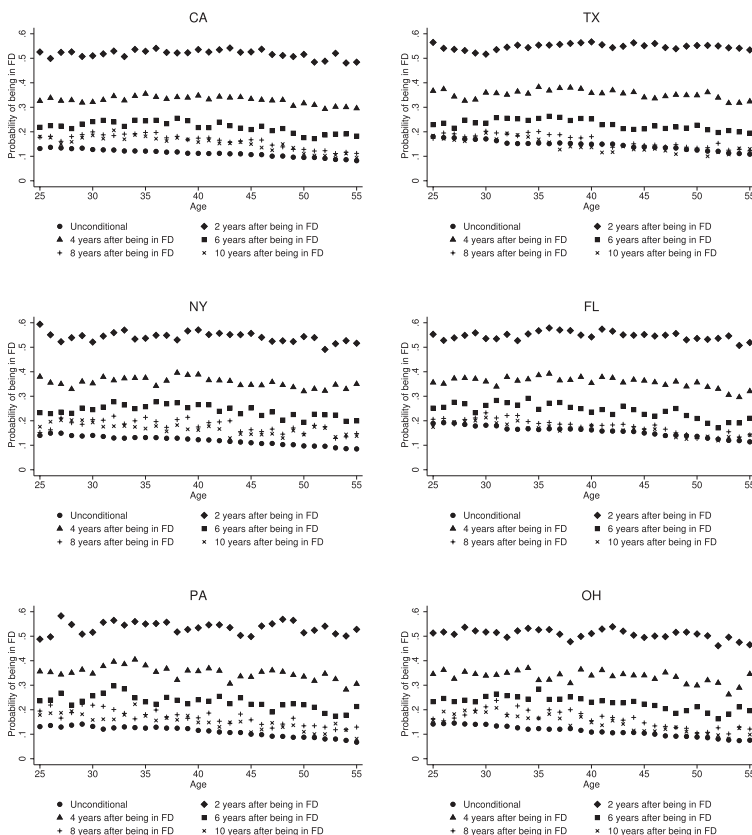


Figure C1
Persistence of financial distress over the life cycle and across states (debt)
 Source: Federal Reserve Bank of New York Consumer Credit Panel/Equifax.

persistence 1–3 years out. Additionally, we still find that discount factor heterogeneity generates long-term persistence of financial distress in models with or without delinquency: All models in Table D2 generate significant persistence 5–10 years out.

Appendix E. Ignoring the Persistence and Concentration of Financial Distress in the Baseline Model

In this section we examine the restrictions that the persistence and concentration of financial distress might impose on the distribution of discount factor heterogeneity. In other words, if we were to estimate models like our benchmark model (with delinquency and bankruptcy) without targeting the persistence or concentration of financial distress, would we recover the same implied distribution of discount factor heterogeneity and would such a model replicate the previously documented patterns of distress? We find that the answer to both of these questions is “no”, suggesting the persistence and concentration of financial distress impose additional restrictions on the distribution of discount factor heterogeneity beyond what the distribution of wealth suggests.

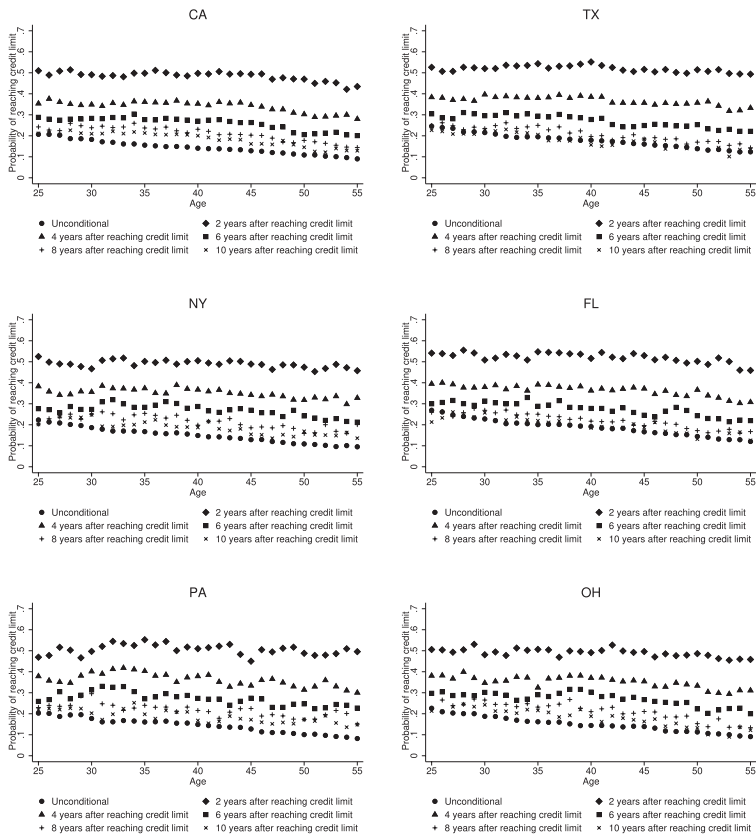


Figure C2
Persistence of financial distress over the life cycle and across states (credit limit)
 Source: Federal Reserve Bank of New York Consumer Credit Panel/Equifax.

To see the answer to the first question, Table E1 presents the results of estimating our benchmark model targeted to match the incidences of financial distress and bankruptcy, the life-cycle profile of wealth-to-earnings ratios, and moments summarizing cross-sectional wealth inequality. The key takeaway from Columns 1 and 2 is that regardless of the income process used, the estimated values for β_L and p_L are very different than those in the main text. For example, the parameter estimates in the main text imply for the RIP specification an average β of 0.92 with a standard deviation of 0.09. Meanwhile, the estimates in Column 1 imply an average β of 0.85 with a standard deviation of 0.2.

Next, the results in Table E2 show that models (with discount factor heterogeneity) that ignore the persistence and concentration of financial distress imply counterfactual patterns for these moments. Both Columns 1 and 2 of this table show that the persistence of financial distress never falls below 0.5 even 10 years out, whereas the data suggest a conditional probability below 0.2 after 10 years. Additionally, both Columns 1 and 2 imply much less concentration of financial distress compared to the data.

Table D1
Parameter values for β -het models when estimating β_H

	DQBK RIP	DQBK HIP	No-DQ RIP	No-DQ HIP
Earnings threshold in DQ τ	68.333 (39.709)	29.830 (8.898)	–	–
Earnings threshold in BK τ	–	–	24.403 (9.227)	12.607 (3.378)
Low discount factor β_L	0.816 (0.041)	0.811 (0.055)	0.713 (0.072)	0.695 (0.066)
High discount factor β_H	1.000 (0.037)	0.970 (0.021)	0.998 (0.044)	1.000 (0.045)
Discharge shock to DQ debt γ	0.146 (0.031)	0.155 (0.020)	–	–
Share of pop. of type L p_L	0.384 (0.111)	0.327 (0.106)	0.356 (0.056)	0.314 (0.061)

Asymptotic standard errors are reported in parentheses.

Table D2
Fit of β -het models with β_H estimated

	Data	DQBK RIP (1)	DQBK HIP (2)	No-DQ RIP (3)	No-DQ HIP (4)
(1) FD rate, ages 25–34 (%)	13.42	19.94	24.29	14.72	15.11
(2) FD rate, ages 35–44 (%)	11.69	12.35	12.25	10.34	9.88
(3) FD rate, ages 45–55 (%)	9.35	9.81	8.36	9.92	8.66
(4) BK rate, ages 25–34 (%)	0.87	1.42	1.55	14.72	15.11
(5) BK rate, ages 35–44 (%)	1.00	0.70	0.72	10.34	9.88
(6) BK rate, ages 45–55 (%)	0.78	0.65	0.57	9.92	8.66
(7) Wealth-to-earnings, ages 25–34	1.12	0.67	0.52	0.67	0.82
(8) Wealth-to-earnings, ages 35–44	2.04	2.26	1.47	2.24	2.69
(9) Wealth-to-earnings, ages 45–55	3.29	4.38	3.12	4.37	5.40
(10) Wealth Gini	0.76	0.69	0.70	0.68	0.66
(11) Wealth ($P75 - P25$)/ $P50$	3.62	5.02	5.67	4.32	4.83
(12) Average Pr(FD+1 FD)	0.69	0.62	0.62	0.00	0.00
(13) Average Pr(FD+3 FD)	0.41	0.37	0.37	0.16	0.16
(14) Average Pr(FD+5 FD)	0.27	0.29	0.28	0.22	0.21
(15) Average Pr(FD+8 FD)	0.16	0.23	0.21	0.25	0.25
(16) Average Pr(FD+10 FD)	0.15	0.20	0.18	0.23	0.22
(17) 70th percentile of FD	0.08	0.05	0.08	0.10	0.09
(18) 80th percentile of FD	0.23	0.20	0.23	0.30	0.27
(19) 90th percentile of FD	0.52	0.49	0.50	0.59	0.56
χ^2		0.81	0.81	0.78	0.88

Appendix F. A Model without DQ

In this section we provide details for the model without DQ mentioned in the main text.

We assume that in each period, households may formally default (i.e., file for bankruptcy) on existing debt. Like in our benchmark model in the main text, households trade off the advantages and disadvantages of formal default. 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. Doing so comes

Table E1
Parameter values for DQBK β -het models without FD targets

	RIP (1)	HIP (2)
Earnings threshold in DQ τ	5.941 (3.636)	7.756 (4.692)
Low discount factor β_L	0.549 (0.086)	0.643 (0.057)
High discount factor β_H	1.000 [†]	1.000 [†]
Discharge shock to DQ debt γ	– 0.026 (0.023)	– 0.021 (0.011)
Share of pop. of type L p_L	0.329 (0.041)	0.266 (0.253)

Asymptotic standard errors are reported in parentheses. [†] denotes a parameter fixed by assumption.

Table E2
Fit of DQBK β -het models with and without FD moments

	Data	No persistence or concentration of FD	
		RIP (1)	HIP (2)
(1) FD rate, ages 25–34 (%)	13.42	17.12	18.05
(2) FD rate, ages 35–44 (%)	11.69	12.37	13.14
(3) FD rate, ages 45–55 (%)	9.35	11.56	10.35
(4) BK rate, ages 25–34 (%)	0.87	1.92	1.77
(5) BK rate, ages 35–44 (%)	1.00	0.89	0.95
(6) BK rate, ages 45–55 (%)	0.78	0.69	0.78
(7) Wealth-to-earnings, ages 25–34	1.12	0.69	0.83
(8) Wealth-to-earnings, ages 35–44	2.04	2.41	2.85
(9) Wealth-to-earnings, ages 45–55	3.29	4.61	5.70
(10) Wealth Gini	0.76	0.68	0.65
(11) Wealth ($P75 - P25$)/ $P50$	3.62	3.61	3.66
(12) Average Pr(FD+1 FD)	0.69	0.80	0.81
(13) Average Pr(FD+3 FD)	0.41	0.70	0.66
(14) Average Pr(FD+5 FD)	0.27	0.67	0.59
(15) Average Pr(FD+8 FD)	0.16	0.65	0.53
(16) Average Pr(FD+10 FD)	0.15	0.64	0.50
(17) 70th percentile of FD	0.08	0.00	0.00
(18) 80th percentile of FD	0.23	0.05	0.05
(19) 90th percentile of FD	0.52	0.35	0.33
χ^2		0.39	0.40

with two disadvantages, however. In the period of bankruptcy, a proportion of income, τ , is lost.²³ Additionally, in that period, consumption equals income—neither saving nor borrowing is allowed. In this environment, lifetime utility can be written as

$$G_{i,n}(z, \varepsilon, a) = \max \left\{ \underbrace{V_{i,n}(z, \varepsilon, a)}_{\text{Pay}}, \underbrace{B_{i,n}(z, \varepsilon)}_{\text{Bankruptcy}} \right\}, \quad (\text{F1})$$

²³ Chatterjee, Corbae, and Rios-Rull (2008) build a model where no punishment is required after filing bankruptcy. There, asymmetric information is crucial in creating incentives for debt repayment because households signal their type by paying back their debt.

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),

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

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

$$B_{i,n}(z, \varepsilon) = u(\min\{y_{i,n}(z, \varepsilon), \tau\}) + \varrho_n \beta \mathbb{E}[G_{i,n+1}(z', \varepsilon', 0) | z]. \tag{F2}$$

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

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

subject to (F3)

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

$$c \geq 0.$$

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.

$$q_{i,n}(z, a') = \frac{1}{1+r} \varrho_n \mathbb{E}[R_{i,n}(z', \varepsilon', a') | z]. \tag{F4}$$

Looking at this equation, it is very clear why the price function (or interest rate) 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 (F1) to (F3) are satisfied and prices satisfy the zero-profit condition (F4).

Appendix G. Calibrating the Non-DQ Model to Data on Bankruptcy

In the main text the No-DQ models are estimated to match financial distress even though these models only allow for formal default (i.e., bankruptcy). This was done for two reasons. First, in reality, debt in delinquency has a very low value indicating that most of the debt in delinquency is discharged at some point. Our interpretation of delinquent debt as equivalent to defaulted debt reflects this observation. Second, our main goal is to analyze the performance of any model in terms of the persistence of financial distress, and, in general, there are no good data on the persistence of bankruptcy.

Nevertheless, in this section we explore what our No-DQ models with no discount factor heterogeneity imply for the persistence of bankruptcy (interpreted as financial distress) when estimated to match bankruptcy rates over the life cycle. Overall, we find that it is difficult to generate much persistence of bankruptcy when there is no delinquency or discount factor heterogeneity.

We consider two estimation procedures. In the first case, we only target bankruptcy rates, following Livshits, MacGee, and Tertilt (2007). In the second case, we target both bankruptcy rates and wealth moments to see if these models can simultaneously generate a low bankruptcy rate over the life cycle in conjunction with reasonable wealth accumulation and dispersion patterns.

Table G1
Parameter values for no-DQ no-het models estimated to BK and wealth targets

Parameter	Only BK moments		BK and wealth moments	
	RIP (1)	HIP (2)	RIP (3)	HIP (4)
Earnings threshold in BK τ	1.101 (0.147)	1.305 (0.147)	9.842 (3.041)	5.739 (0.482)
Discount factor β	0.771 (0.120)	0.742 (0.113)	0.953 (0.014)	0.955 (0.024)

Asymptotic standard errors are reported in parentheses.

Table G2
Fit of no-DQ no-het models estimated to BK and wealth targets

	Data	Only BK moments		BK and wealth moments	
		RIP (1)	HIP (2)	RIP (3)	HIP (4)
(1) BK rate, ages 25–34 (%)	0.87	1.10	0.97	13.12	14.72
(2) BK rate, ages 35–44 (%)	1.00	0.81	0.73	1.41	10.34
(3) BK rate, ages 45–55 (%)	0.78	0.82	1.02	0.03	9.92
(4) Wealth-to-earnings, ages 25–34	1.12	-1.02	-1.52	0.37	0.67
(5) Wealth-to-earnings, ages 35–44	2.04	-1.53	-1.89	0.74	2.24
(6) Wealth-to-earnings, ages 45–55	3.29	-1.26	-1.24	2.18	4.37
(7) Wealth Gini	0.76	-0.46	-0.48	0.63	0.68
(8) Wealth ($P75 - P25$)/ $P50$	3.62	-1.41	-3.02	2.57	4.32
(9) Average Pr(FD+1 FD)	0.69	0.00	0.00	0.00	0.00
(10) Average Pr(FD+3 FD)	0.41	0.03	0.02	0.04	0.16
(11) Average Pr(FD+5 FD)	0.27	0.03	0.03	0.02	0.22
(12) Average Pr(FD+8 FD)	0.16	0.04	0.03	0.01	0.25
(13) Average Pr(FD+10 FD)	0.15	0.04	0.03	0.01	0.23
(14) 70th percentile of FD	0.08	0.00	0.00	0.00	0.10
(15) 80th percentile of FD	0.23	0.00	0.00	0.12	0.30
(16) 90th percentile of FD	0.52	0.19	0.31	0.40	0.59
χ^2		0.08	0.13	3.33	1.65

Table G1 presents the estimation results under both sets of targets. Columns 1 and 2 of this table show that for these models to generate a low incidence of bankruptcy over the life cycle, a very harsh default punishment is required. Note that the estimated values for τ , the earnings threshold when in bankruptcy, are much lower than any of the estimates in the main text. Next, Columns 3 and 4 show that matching wealth moments not only require higher discount factors (i.e., more patience) but also a less severe bankruptcy penalty in order to generate any default.

Next, Table G2 presents the empirical performance of these models. As can be seen from Columns 1 and 2, models that are only targeted to match bankruptcy rates generate counterfactual wealth accumulation patterns. This occurs because most debt is essentially riskless since bankruptcy is rare. In contrast, Columns 3 and 4 reveal that once wealth moments are added to the estimation, these models imply excessively high bankruptcy rates, particularly early in the life cycle. Importantly, though, all of these models, except for Column 4, generate very little persistence of financial distress. Only the HIP specification targeted to match bankruptcy and wealth moments (Column 4) generates some persistence of financial distress at long horizons (5+ years). However, note that this model implies bankruptcy rates roughly 10 times larger than what is found in the data.

References

- Altonji, J. G., and L. M. Segal. 1996. Small-sample bias in GMM estimation of covariance structures. *Journal of Business and Economic Statistics* 14:353–66.
- Ampudia, M., H. van Vlokhoven, and Z. Dawid. 2016. Financial fragility of euro area households. *Journal of Financial Stability* 27:250–62.
- Athreya, K. 2008. Default, insurance, and debt over the life-cycle. *Journal of Monetary Economics* 55:752–74.
- Athreya, K., J. M. Sánchez, X. S. Tam, and E. R. Young. 2015. Labor market upheaval, default regulations, and consumer debt. *Review of Economic Dynamics* 18:32–52.
- . 2018. Bankruptcy and delinquency in a model of unsecured debt. *International Economic Review* 59:593–623.
- Athreya, K., X. Tam, and E. Young. 2009. Unsecured credit markets are not insurance markets. *Journal of Monetary Economics* 56:83–103.
- Becker, S. G., and C. B. Mulligan. 1997. The endogenous determination of time preference. *Quarterly Journal of Economics* 112:729–58.
- Belzil, C., and J. Hansen. 1999. Subjective discount rates, intergenerational transfers and the return to schooling. Technical Report, IZA Discussion paper series, no. 60.
- Blundell, R., L. Pistaferri, and I. Preston. 2008. Consumption inequality and partial insurance. *American Economic Review* 98:1887–921.
- Bowsher, C. 2002. On testing overidentifying restrictions in dynamic panel data models. *Economic Letters* 77:221–20.
- Brunetti, M., G. Elena, and T. Costanza. 2016. Is financial fragility a matter of illiquidity? an appraisal for Italian households. *Review of Income and Wealth* 62:628–49.
- Cagetti, M. 2003. Wealth accumulation over the life cycle and precautionary savings. *Journal of Business and Economic Statistics* 28:339–53.
- Chamberlain, G. 1982. Multivariate regression models for panel data. *Journal of Econometrics* 18:5–46.
- . 1984. Panel data. In *Handbook of econometrics*, eds. Z. Griliches and M. D. Intriligator, vol. 2, pp. 1247–318. Amsterdam: North-Holland Press.
- Chatterjee, S., D. Corbae, M. Nakajima, and J.-V. Ríos-Rull. 2007. A quantitative theory of unsecured consumer credit with risk of default. *Econometrica* 75:1525–91.
- Chatterjee, S., D. Corbae, and J. V. Ríos-Rull. 2008. A finite-life private-information theory of unsecured consumer debt. *Journal of Economic Theory* 142:149–77.
- De Nardi, M., and G. Fella. 2017. Saving and wealth inequality. *Review of Economic Dynamics* 26:280–300.
- De Nardi, M., E. French, and J. B. Jones. 2016. Medicaid Insurance in Old Age. *American Economic Review* 106:3480–520.
- Favilukis, J., S. C. Ludvigson, and S. V. Nieuwerburgh. 2017. The macroeconomic effects of housing wealth, housing finance, and limited risk sharing in general equilibrium. *Journal of Political Economy* 125:140–223.
- Frederick, S., G. Loewenstein, and T. O'Donoghue. 2002. Time discounting and time preference: A critical review. *Journal of Economic Literature* 40:351–401.
- Fulford, S. L., and S. Schuh. 2017. Credit card utilization and consumption over the life cycle and business cycle. Technical Report, Federal Reserve Bank of Boston.
- Gorbachev, O., and M. J. Luengo-Prado. 2016. The credit card debt puzzle: The role of preferences, credit risk, and financial literacy. Technical Report, Federal Reserve Bank of Boston.

- Gross, D. B., and N. S. Souleles. 2002. Do liquidity constraints and interest rates matter for consumer behavior? evidence from credit card data. *Quarterly Journal of Economics* 117:149–85.
- Guvenen, F. 2009. An empirical investigation of labor income processes. *Review of Economic Dynamics* 12:58–79.
- Hall, Robert, E., and F. Mishkin. 1982. The sensitivity of consumption to transitory income: Estimates from panel data on households. *Econometrica* 50:461–81.
- Hatchondo, J. C., L. Martinez, and J. M. Sánchez. 2015. Mortgage defaults. *Journal of Monetary Economics* 76:173–90.
- Hausman, J. A. 1979. Individual discount rates and the purchase and utilization of energy-using durables. *Bell Journal of Economics* 10:33–54.
- Heathcote, J., K. Storesletten, and G. L. Violante. 2014. Consumption and labor supply with partial insurance: An analytical framework. *American Economic Review* 104:2075–126.
- Jappelli, T., and M. Pagano. 1999. The welfare effects of liquidity constraints. *Oxford Economic Papers* 51:410–30.
- Jappelli, T., M. Pagano, and M. Di Maggio. 2013. Households? indebtedness and financial fragility. *Journal of Financial Management Markets and Institutions* 1:26–35.
- Kaplan, G., and G. L. Violante. 2010. How much consumption insurance beyond self-insurance? *American Economic Journal: Macroeconomics* 2:53–87.
- Krusell, P., and A. Smith. 1998. Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy* 106:867–96.
- . 1999. On the welfare effects of eliminating business cycles. *Review of Economic Dynamics* 2:245–72.
- Lawrance, E. C. 1991. Poverty and the rate of time preference: Evidence from panel data. *Journal of Political Economy* 99:54–77.
- Livshits, I., J. MacGee, and M. Tertilt. 2007. Consumer bankruptcy: A fresh start. *American Economic Review* 97:402–18.
- Lusardi, A. 2011. Americans' financial capability. Working Paper.
- Lusardi, A., D. Schneider, and P. Tufano. 2011. Financially fragile households: Evidence and implications. *Brookings Papers on Economic Activity*, 2011:83–134.
- Meier, S., and C. Sprenger. 2017. Present-biased preferences and credit card borrowing. *American Economic Journal: Applied Economics* 2:193–210.
- Mustre-del Río, J. 2015. Wealth and labor supply heterogeneity. *Review of Economic Dynamics* 18:619–34.
- Parker, J. A. 2017. Why don't households smooth consumption? Evidence from a \$25 million experiment. *American Economic Journal: Macroeconomics* 9:153–83.
- Picchi, A. 2015. The financial fragility of the american household. *CBSNews*, May 29. <https://www.cbsnews.com/news/the-financial-fragility-of-the-american-household/>.
- Roodman, D. M. 2009. A note on the theme of too many instruments. *Oxford Bulletin of Economics and Statistics* 71:135–58.
- Samwick, A. A. 1998. Discount rate heterogeneity and social security reform. *Journal of Development Economics* 57:117–46.
- Warner, J. T., and S. Pleeter. 2001. The personal discount rate: Evidence from military downsizing programs. *American Economic Review* 91:33–53.
- Zeldes, S. 1989. Consumption and liquidity constraints: An empirical investigation. *Journal of Political Economy* 97:305–46.