ECONOMIC RESEARCH
FEDERAL RESERVE BANK OF ST. LOUIS

WORKING PAPER SERIES

The Effects of Macroeconomic Shocks: Household Financial Distress Matters

Authors	Kartik Athreya, Ryan Mather, José Mustre-del-Río, and Juan M. Sánchez
Working Paper Number	2019-025G
Revision Date	September 2023
Citable Link	https://doi.org/10.20955/wp.2019.025
Suggested Citation	Athreya, K., Mather, R., Mustre-del-Rio, J., Sanchez, J.M., 2023; The Effects of Macroeconomic Shocks: Household Financial Distress Matters, Federal Reserve Bank of St. Louis Working Paper 2019-025. URL https://doi.org/10.20955/wp.2019.025

Federal Reserve Bank of St. Louis, Research Division, P.O. Box 442, St. Louis, MO 63166

The Effects of Macroeconomic Shocks: Household Financial Distress Matters

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

September 11, 2023

Abstract

When a macroeconomic shock arrives, variation in household balance-sheet health (captured by the presence of financial distress "FD"), leads to differential access to credit, and hence a distribution of consumption responses. As we document, though, over the past two recessions, households in prior FD also experienced macroeconomic shocks more intensely than others, leading to a distribution of shock severity. Quantifying the importance of each dimension of heterogeneity (FD or shock severity) for consumption requires a structural model. We find that heterogeneity in FD matters more than dispersion in shock severity for shaping the responses of individual and aggregate consumption to any shock.

Keywords: Consumption, Credit Card Debt, Recession, Bankruptcy, Foreclosure, Mortgage, Delinquency, Financial Distress, Inequality, Poverty.

JEL Codes: D31, D58, E21, E44, G11, G12, G21.

*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.
‡Department of Economics, University of Wisconsin-Madison; e-mail: mather2@wisc.edu.
§Federal Reserve Bank of Richmond; e-mail: kartik.athreya@rich.frb.org

1 Introduction

Many recent studies have shown that accurately capturing heterogeneity in household balance sheets is crucial for modeling the aggregate and distributional consequences of shocks.[1] There are two potential reasons for this: First, it could be that a shock is not experienced uniformly, and different types of households receive different "shares" of aggregate losses. Second, it could be that households are different from one another in persistent ways, and therefore, would respond differently even if they had all experienced the same shock.

In this paper, we evaluate how these two dimensions of heterogeneity matter in shaping consumption responses to aggregate shocks. Our measure of household heterogeneity is based on their financial vulnerability, or what we call financial distress (FD). FD captures whether households are over 30 days delinquent on paying back unsecured debt. While somewhat non-standard, we think this is a useful measure of financial vulnerability as it is easily observed in credit bureau data, is very persistent at the individual level, and projects well on household-level marginal propensities to consume (MPCs) in response to shocks.[2] Thus, heterogeneity in FD might be a reason why households respond differently to the same shock. We also show that during the last two recessions, the burden of aggregate shocks was worse among households who were in greater FD before each recession. Hence, FD might also be capturing heterogeneity in aggregate shock exposure.

Gauging the effects on consumption of differences in FD across households requires a departure from the standard incomplete market model.[3] We build a structural model that incorporates housing, mortgages, and unsecured debt with both formal default (bankruptcy) and informal default via non-repayment (delinquency). Informal default is a particularly important feature as it helps formalize FD in our model. We structurally estimate critical parameters of the model to match key moments including the household-level persistence of FD observed in the data. This exercise suggests that matching the distribution of FD, and its persistence, in and of itself implies a significant degree of ex-ante heterogeneity in the population.

[1] See Kaplan and Violante (2014) and Aruoba, Elul, and Kalemli-Ozcan (2018).

[2] These points are discussed with more detail in Section 2.1.

[3] Although the standard model was developed by Huggett (1993) and Aiyagari (1994), to be more precise we are referring here to the more recent quantitative versions in Berger, Guerrieri, Lorenzoni, and Vavra (2018) and Kaplan, Mitman, and Violante (2020) that allow for housing and mortgages.

Next, to help assess the importance of heterogeneity in shock exposure we construct aggregate shocks that mimic the observed relationship between FD and shocks observed in the last two recessions. Our stylized version of the Great Recession (GR) is characterized by a decline in house prices that is more severe among households in greater FD. Similarly, our stylized version of the COVID-19 pandemic (CV19) is characterized by labor-income losses that are also more severe among households in greater FD. With the model and the shocks as described, we have a credible laboratory to gauge the importance of differences in shocks and differences in people for shaping consumption responses to aggregate shocks.

To assess the importance of heterogeneity in households versus shock exposure, we examine the response of consumption both at the macro- and micro-level. As is common practice, we largely focus on the response of aggregate consumption to both of the shocks we model. However, because aggregate responses can mask significant differences at the micro-level, we additionally consider how consumption inequality and consumption-based poverty respond.[4] The former gives us a sense of changes in the cross-section, while the latter gives us a sense of changes amongst the most disadvantaged.

We find that heterogeneity across households in FD is important as it *alters* (i.e. either amplifies or attenuates) the response of the aforementioned measures by nearly 25 percent, on average, depending on the type of shock considered. House price declines, similar to those experienced during the GR, tend to reduce aggregate consumption, but also consumption inequality and poverty. All these effects are about 30 percent larger compared to an alternative version of our model that excludes FD. Labor income declines, similar to those experienced during the CV19, reduce aggregate consumption but increase consumption inequality and poverty. Allowing for FD amplifies the drop in aggregate consumption, but attenuates the increase in inequality and poverty. In an absolute sense, FD changes these responses by about 20 percent, on average.

In contrast, we find a smaller quantitative role for heterogeneity in shock exposure. We find fairly similar responses of aggregate consumption, inequality, and poverty in counterfactual economies where aggregate shocks are uncorrelated with prior FD. This suggests much of the previously mentioned quantitative results are not simply due to similar households responding to different shocks. Rather, our model-based analysis

[4]Consumption-based poverty is the proportion of the population that consumes below the cost of basic needs. See Armstrong et al. (2022).

suggests households, that differ amongst other things in their FD status, would respond differently to the same shocks.

Beyond consumption, we also find that FD has important consequences for the housing market. Measures like housing leverage and mortgage default rates are notably higher in our baseline model with FD compared to a counterfactual economy that excludes it. This difference arises because in a model with FD, borrowing rates account for the fact that homeowners can extract equity from their houses or default on mortgage payments in order to pay unsecured debt. As a result, in a model with FD, borrowing rates are more favorable for homeowners. This encourages homeownership even among riskier types who are more likely to default on mortgage payments.

Overall, our work suggests examining household FD is important for at least two reasons. First, from a practical standpoint, FD is an empirically tractable and valuable "tagging" mechanism. In Section 5 we show within our model that individuals in FD have higher MPCs (out of house-price and income shocks) than individuals not in FD. Moreover, differences in MPCs by FD are very similar to differences obtained when sorting individuals by preference type, which is unobservable and the primitive driving the differences in MPCs. Thus, FD not only allows us to identify high versus low MPC individuals, it also allows us to "back out" the underlying (and unobservable) preference type.[5] Second, the results we present in Section 6 show that modeling FD in and of itself matters for understanding aggregate consumption, consumption inequality and poverty, and the housing market. Holding all else constant, the predictions of our model with FD differ from those of a model that allows for borrowing up to an ad hoc limit. This suggests that the state- and type-specific borrowing constraint that our model with FD delivers differs in a quantitatively meaningful way compared to the ad hoc borrowing constraint more commonly used in the literature. However, as we study two types of shocks and various measures our paper effectively provides some guidance for when the tractability of a model with an ad hoc borrowing constraint outweighs the additional realism of a

[5]While our analysis leverages discount factor heterogeneity, this is still a stand-in for a variety of other unobserved demands for consumption within the household arising from a variety of sources. For example, Becker and Mulligan (1997) shows how addictions, uncertainty, and other variables affect the degree of time preference. Thus, 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.

model with FD, and vice versa.

The remainder of the paper is structured as follows. Below we provide a brief literature review and motivating evidence that FD is a valuable measure of household/consumer vulnerability, including in the context of the broader literature. Section 2 provides further details on the empirical relationship between FD and aggregate shocks during the past two recessions. Section 3 develops our model of consumption, debt, and default. Section 4 addresses the model parameterization and estimation, along with the details of calibration of the aggregate shocks. Section 5 provides validation of the model against external information on the responsiveness of consumption to shocks. It also provides evidence of the usefulness of FD as a "tagging" mechanism. Section 6 contains our main quantitative results showing the importance of FD alone vis-á-vis its correlation with aggregate shocks. Finally, Section 7 concludes.

1.1 Related literature

Given our interest in how FD affects the transmission of shocks into consumption, our paper is strongly tied to several research strands in macroeconomics, which we discuss below.

Is FD a new theoretical concept that has not been captured in the models the literature used so far? The concept of FD is relatively new and its importance for the response to macro shocks has not been studied. There is literature in macroeconomics studying the related concept of household bankruptcy that started in with Athreya (2002), Chatterjee, Corbae, Nakajima, and Ríos-Rull (2007), and Livshits, MacGee, and Tertilt (2007). In those models, bankruptcy means that debt is forgiven and that the household cannot file bankruptcy again for seven years. This concept is a very extreme form of financial distress that is much less common than the type we focus on, where people simply go delinquent on debt payments but do not formally default. While less than 0.5 percent of households file bankruptcy in a given year, we find that about 15 percent are in FD according to our definition. In other words, around 30 times as many people are affected by FD in a year as experience bankruptcy. In previous work, we have developed the theoretical concept of financial distress, but this was done at the micro-economic level. We studied something related to financial distress, referred to as informal default, in two previous papers: Athreya, Sánchez, Tam, and Young (2015) and Athreya, Sánchez, Tam, and Young (2017). As in the current model, during informal default households can skip payments and they are charged a penalty rate for the next period. The concept of FD specifically was introduced in a recent publication, Athreya, Mustre-del Río, and Sánchez (2019). The main contribution of that paper was to demonstrate that discount-factor heterogeneity allows the model to reproduce the persistence of FD.[6]

Our examination of the non-repayment of debt and its importance for understanding the response of consumption to house-price shocks is closely related to recent work aimed at understanding dynamics in the wake of house price movements. However, our work differs from this strand of the literature because we incorporate formal and informal default as alternative margins of adjustment in the financial asset market. Berger, Guerrieri, Lorenzoni, and Vavra (2018) was the first paper to study how prices affect consumption in a quantitative heterogeneous-agent model with incomplete markets and liquidity constraints. They show how consumption responses depend on factors such as the level and distribution of debt, the size and history of house price shocks, and the level of credit supply. The idea of incorporating mortgage default in a model with exogenous house price shocks follows Corbae and Quintin (2015) and Hatchondo, Martinez, and Sánchez (2015). In that regard, our work is related but different than papers with similar life-cycle models but trying to account for the joint evolution of house prices and consumption during the GR (Garriga and Hedlund, 2017; Kaplan, Mitman, and Violante, 2020).

Our results on house-price shocks are also related to the empirical work of Aruoba, Elul, and Kalemli-Ozcan (2018). They decompose the effect of declining house prices on consumption into a wealth effect, household financial constraints, and bank health. Critically, they find little evidence of a wealth effect, yet about 40-45 percent of the response of consumption can be explained by tightening household financial constraints. Our model decomposition suggests some of the effects of FD operate through the structure it imposes on debt holdings and the price of debt across households. Indeed, an alternative model with a fixed borrowing constraint (which precludes the discussion of tightening financial constraints) does not generate the same responses of consumption to house

[6] In that paper, we also showed that without differences in discount factors is impossible to capture the persistence of financial distress. Basically, we need some agents who consistently are willing to borrow more (to obtain a bit more resources today) in exchange for a higher risk of getting into FD in the next period.

price shocks that our baseline model does. In this sense, our model is consistent with the view that a fraction of the response of consumption to house price shocks is due to tightening credit constraints.

The analysis of how FD affects the transmission of income shocks into consumption is related to a set of papers that emphasize the modeling of delinquency or bankruptcy and how it shapes macroeconomic fluctuations. The main difference between that literature and our work is that we consider other channels by which delinquency or bankruptcy shape aggregate responses. For example, while Herkenhoff and Ohanian (2012) and Herkenhoff (2013) emphasize the importance of default for the dynamics of unemployment, Auclert and Mitman (2019) examine how the default choice is amplified through the Keynesian channels of aggregate demand (via sticky prices and aggregate demands externalities). Viewed through the lens of our model, those papers focus on how FD as an alternative margin of adjustment affects subsequent macroeconomic outcomes. Our contribution is to also analyze how FD matters through the ex-ante heterogeneity it encodes and its correlation with aggregate shocks.

Next, our conclusion that dispersion in household consumption responses (an endogenous outcome) is mostly due to heterogeneity in FD (or, as we show in Section 5 heterogeneity in MPCs) across households rather than heterogeneity in the shocks they receive is akin to the work of Berger and Vavra (2019) on pricing behavior at the firm level. They document that item-level price change dispersion is both countercyclical and highly correlated with exchange rate pass-through. Using a workhorse open-economy model they find these facts support an important role for time-varying responsiveness, whereas time-varying shock volatility is less important. Our results suggest that household-level differences in responsiveness (driven by differences in FD) are more important for shaping the distribution and aggregate level of consumption than differences in the shocks these households receive.

The approach of using information from households in the left-tail of the wealth distribution to identify heterogeneity is related but different than previous work that has mostly used the right-tail of the wealth distribution (Krusell and Smith., 1998). In this sense, our findings are in line with Parker (2017), who notes how a "main finding is that the majority of lack of consumption smoothing is predicted by a simple measure that can be interpreted as impatience."

Our estimation procedure is also related to several papers using individual-level data and structural models to identify preference heterogeneity more generally. Aguiar, Bils, and Boar (2020) find that both discount factor heterogeneity and heterogeneity in the intertemporal elasticity of substitution (IES) are necessary to generate the correct individual consumption responses to income shocks. Similarly, Calvet, Campbell, Gomes, and Sodini (2019) also finds support for heterogeneity in discount factors and the IES when looking at spending and savings patterns from Swedish households. Mustre-del Río (2015) finds that substantial dispersion in the disutility of work is needed to match dispersion in labor supply across individuals that cannot be explained by wage differences alone. Finally, Gregory, Menzio, and Wiczer (2021) also find evidence of substantial heterogeneity across workers using data from the Longitudinal Employer-Household Dynamics (LEHD) dataset. Compared to those papers we show how data on FD and homeownership identify a correlation between discount factors and preference for homeownership that shapes the predictions of poverty in response to house-price and income shocks.

2 Empirical evidence

2.1 Why FD as a measure of vulnerability

In this section we motivate household FD as a useful and timely measure of financial vulnerability. We define FD as a case when an individual has a credit card account at least thirty days delinquent at some point during the year (i.e. *DQ30*). We also present some results for an alternative definition of FD, *CL80*, which is a case when an individual has reached at least 80% of their credit limit over the same time interval.[7] As seen below, our main empirical results are robust to either definition of FD. However, in the quantitative analysis of the subsequent sections we focus on the DQ30 version as it is mostly easily defined in our model.

Either of these definitions of FD is easily measured, timely, and encompassing. They are easily measured because they are built with the New York Fed Consumer Credit Panel (NY FED-CCP), which contains credit reports for millions of Americans. This also makes

[7] We also use some additional metrics for FD in our robustness checks for the relationship between FD and marginal propensities to consume. These are discussed in Appendix A.4.

them timely, because the NY FED-CCP updates the information needed to construct them quarterly, and releases it a few days after the end of the quarter. These variables are encompassing because, unlike other measures, neither requires knowledge of the items on an individual's balance sheets or of the prices needed to compute measures such as net worth or leverage. Moreover, even a near-perfect knowledge of household or individual net wealth may not accurately represent vulnerability. For example, individuals with low levels of net worth may not be constrained.[8] By contrast, seeing an individual become significantly delinquent, or utilizing most if not all unsecured credit, is more telling. Given the costs associated with these actions, it is unlikely that the individuals who take them are unconstrained.

FD is also a valuable measure to look at because it has a relatively high incidence, and is very persistent over the life cycle. Figure 1 is a modified version of a figure from Athreya et al. (2019), updated with our new DQ30 definition of FD. The blue dots along the bottom of the figure show that regardless of age, around 10 to 20% of all individuals find themselves in FD. For those who are in FD, the other markers reveal that this condition is very persistent. For example, the green triangles show that individuals who are in FD today have around a 40% probability of being in FD four years from now. Comparing this to the unconditional average shown by the green dots, being in FD today roughly doubles your odds of being in FD in four years.

Aggregating our measure of FD to the zip-code level, there is strong evidence that being in FD increases vulnerability to macroeconomic shocks. Figure 2 reveals an increasing relationship between zip code-level FD (again using the DQ30 measure) and MPCs. The MPCs plotted in this figure are out of housing shocks and are calculated similarly to Mian, Rao, and Sufi (2013) and Kaplan, Mitman, and Violante (2016) using new auto registrations as the measure of consumption (also at the zip code level).[9] For ease of exposition, we present the average MPCs for different quintiles of FD, ranging from the lowest FD (Q1) to the highest FD (Q5).

As can be seen by looking at the bars in Figure 2, the MPC out of housing shocks

[8] Think of those in middle age who are beginning wealth accumulation for retirement or those financially assisted by relatives. At the other end of the spectrum, those with high "observable" wealth or net worth may be significantly constrained due to debt and other potentially more informal future obligations not easily seen or consumption commitments.

[9] See Appendix A.4 for details and robustness of this relationship.

Figure 1: The Incidence of Persistence of FD Over the Life Cycle

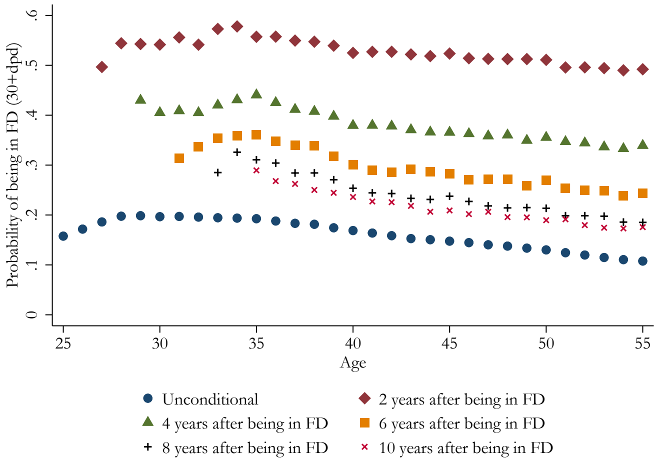

Source: This figure plots the average probability of being in FD, defined as an individual having a credit card account 30 days or more delinquent at some point during the year.

increases from less than 1 cent to over 2 cents between the least and most distressed households. For reference, the horizontal line represents the MPC estimated by Mian, Rao, and Sufi (2013). In general, our estimates are slightly smaller than theirs but within the range of estimates reported in Dupor, Mehkari, Li, and Tsai (2019). Significantly, the other bars in Figure 2 show that this finding holds even when we control for housing leverage or measures of income volatility and some local industry shares.[10] That the result survives the inclusion of these terms suggests that FD is capturing something different than conventional measures of vulnerability. While debt levels, income volatility, and the influence of local industries may affect a household's need to go delinquent on debt or use up available short-term credit, these factors cannot explain away the influence of FD.[11] Intuitively, FD status at any given time encodes information about past debt (non)repayment decisions, something not directly captured by current debt or leverage. In this sense, FD may help identify households' attitudes toward debt and repayment, which are crucial to determining the consumption response to shocks.

[10]Specifically, we control for the shares of employment in manufacturing and in services for 2005. More details are present in Appendix A.4.

[11]Relatedly, using a proxy of FD in the Survey of Consumer Finances (SCF), we found that standard demographic characteristics like sex, race, and education explain a very minor (less than 10 percent) portion of FD in the cross-section. Similarly, measures of financial literacy constructed from the survey do not explain FD either. Results are available upon request.

Figure 2: Marginal Propensity to Consume Out of a Dollar Change in Home Prices by Quintile of DQ30 in 2002

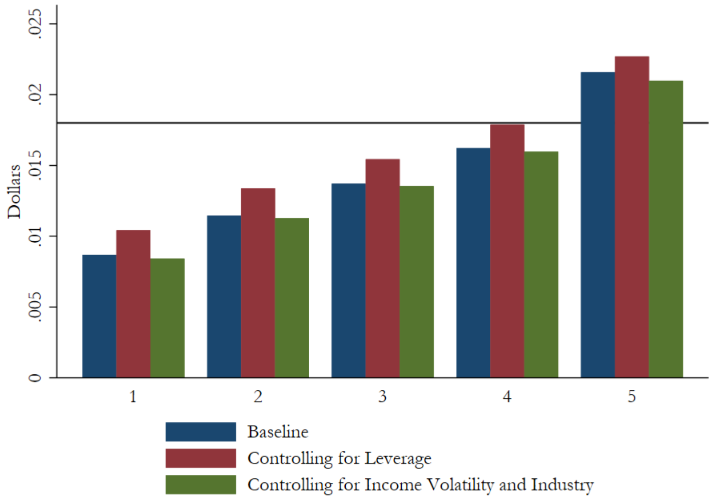

Sources: IRS Survey of Income, FRBNY Consumer Credit Panel/Equifax, Census Bureau, Zillow, Survey of Consumer Finances. Notes: Group means are weighted by the number of owner-occupied housing units per county as of 2006. The horizontal line corresponds to the mean MPC out of autos estimated at the zip code level by Mian, Rao, and Sufi (2013) in their fifth column of Table V.

2.2 FD and its correlation with the size of shocks

Having defined FD and shown its usefulness as an individual measure of vulnerability, this section documents the correlation between FD and aggregate shocks over the past two recessions. Unfortunately, there is no single data source for individual-level data on FD, employment (or income), and wealth. We circumvent this issue by aggregating our individual-level data on FD to the zip-code or county level and merging it with other data sources aggregated to that same level. This allows us to establish two key empirical findings: (i) higher FD before the GR was associated with subsequently larger house-price declines; and (ii) FD before the CV19 pandemic was associated with more significant earnings losses during it. Overall, this suggests that beyond being a relevant measure of individual vulnerability, the distribution of FD across the United States may help us better understand the aggregate and distributional consequences of the past two recessions because of its relationship between the aggregate shocks that precipitated these downturns.

Starting with the GR, the left panel of Figure 3 shows that home values during this event declined the most in higher FD communities. By 2012, regardless of FD, median home prices declined on average by around 15% relative to their 2006 levels. However, home price declines in zip codes with higher FD were twice that, or worse in many cases.

Figure 3: The Correlation Between FD and Aggregate Shocks

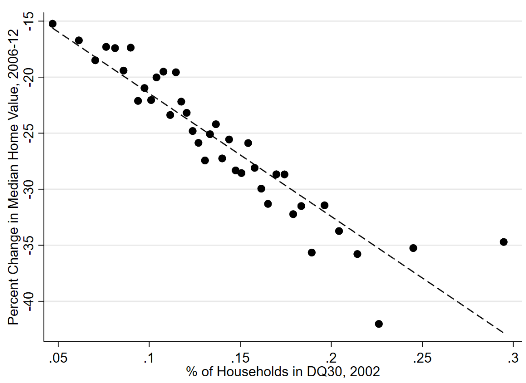

(a) FD and House-price Shocks During GR

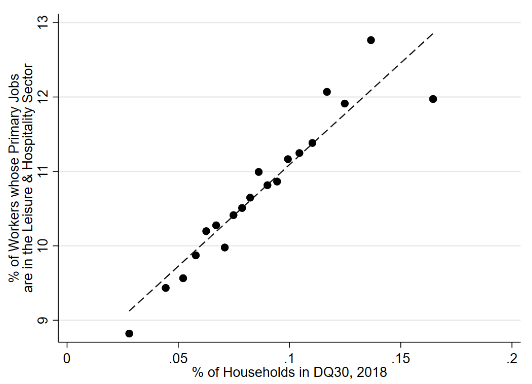

(b) FD and Contact-sensitive Employment

Sources: Census LODES and FRBNY Consumer Credit Panel/Equifax.
Notes: FD is measured as DQ30, which is the share of individuals who are at least thirty days delinquent on a credit card at some point in a given year. For ease of viewing, the data have been divided into forty bins with respect to DQ30, and each dot represents the mean of that bin. In panel (a) each bin is weighted by the housing wealth in each zip code in that bin as of 2006. In panel (b) each bin is weighted by the number of households in each zip code included in the bin.

Perhaps worst of all, households hardest hit were not diversified. Specifically, we find that households with high FD also tended to hold a larger share of their net wealth in their homes. This result implies that when losses are measured as a percentage of net wealth, home value losses are more strongly correlated with FD. In other words, the skewed distribution of home-price losses generated an even more heavily skewed distribution of net wealth losses for regions in higher FD. Appendix Section A.3 illustrates this relationship.

Much like during the GR, the economic consequences of the CV19 pandemic also appear to be correlated with FD. Some suggestive evidence is provided in the right panel of Figure 3. This figure shows a strong and consistently positive relationship between FD incidence at the zip code level (measured by the incidence of DQ30 in 2018) and the share of those areas workers employed in leisure and hospitality. A natural conjecture is that income losses among high FD areas may have been more significant than those in low FD areas.

Survey evidence from Bick and Blandin (2021) does suggest that individuals in higher

FD areas have been more adversely impacted during the CV19 pandemic.[12] Combining our measures of FD at the zip code level with survey responses from Bick and Blandin (2021), we calculate the shares of individuals reporting: (i) no earnings losses (or some increase); and (ii) earnings losses of 50% or more, both relative to earnings in February 2020 (if employed).

Figure 4: Change in Earnings in 2020 by Quintile of FD

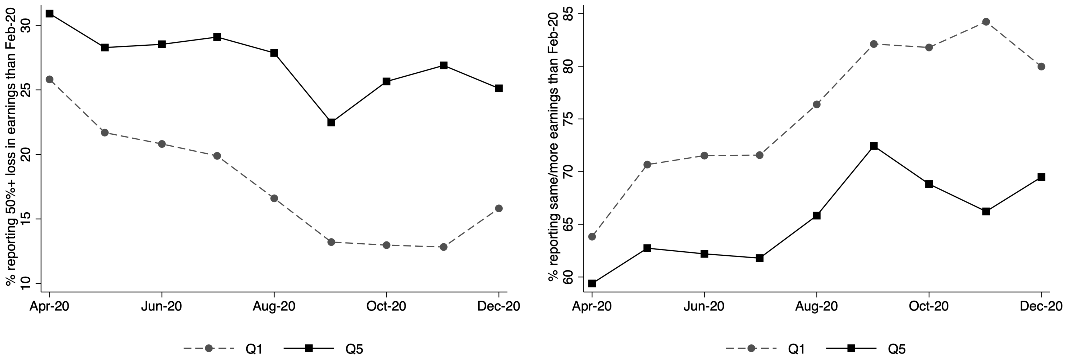

Source: Bick and Blandin (2021) and FRBNY Consumer Credit Panel/Equifax.

The left panel of Figure 4 shows that throughout 2020, individuals living in the most distressed zip codes were consistently more likely to report significant earnings losses compared to individuals living in the least distressed zip codes. Again for expositional simplicity, we group individual responses based on the incidence of FD at the zip code level and focus on differences between individuals living in zip codes with the highest (Q5) and lowest (Q1) incidence of FD.[13] As of December 2020, about 25% of individuals in the highest quintile of FD (Q5) reported earnings losses of at least 50%. In contrast, the comparable figure for individuals in the lowest FD quintile (Q1) is 15%.

This gap in reporting severe earnings losses between Q1 and Q5 is entirely reflected in the incidence of reporting no earnings losses (or increases). As can be seen in the right panel of Figure 4, individuals in Q5 have systematically been less likely to report earnings staying the same or increasing. As of December 2020, about 70% of individuals in Q5 report earnings staying the same. In contrast, around 80% of individuals in Q1 reported their earnings staying the same. Overall, these findings suggest that whether looking at employment in contact-sensitive sectors or actual reported losses, the economic burden of the CV19 pandemic appears to have fallen strongest on the most financially vulnerable.

[12] We are highly appreciative of Alexander Bick and Adam Blandin for sharing their data with us.

[13] Graphs with all quintiles are available upon request.

3 A life-cycle model of housing and FD

As alluded to in the previous section, FD alone may affect the transmission of aggregate shocks into consumption as FD reflects differential access to credit, which leads to differential consumption responses (differences in MPCs). Alternatively, FD may shape the response of consumption somewhat mechanically, because as we documented, prior FD was correlated with the severity of aggregate shocks in the previous two recessions. Given that FD is at least partially endogenous, quantifying these two channels requires a model of debt acquisition, debt repayment, and consumption decisions.

3.1 Agents, markets, and debt default

There is a continuum of finitely lived individuals who are risk averse and discount the future exponentially. All individuals face risk of death in each period and survive to the next period with probability ρ_t, where t denotes age. Agents work for a finite number of periods, retire at age W, and die with certainty at age T (conditional on reaching this terminal age). In what follows, n denotes periods left until the last period of life T, and is naturally related to age by the relation $n = T - t$.

All agents are subject to risk in their income y (specified below). Additionally, agents are allowed to differ in the rate at which they discount the future. Specifically, a share p_L of the population has a discount factor of β_L, while the remaining share has a discount factor of $\beta_H \geq \beta_L$.[14]

Concerning markets, individuals have (limited) access to credit and each period choose nondurable consumption c, housing h, mortgages m', and financial assets (or debt) a'. They may choose to obtain housing services through homeownership or by renting.

Agents enter each period either as nonhomeowners or homeowners. Rental houses are of size h^R, while owner-occupied houses vary in discrete sizes $h' \in \{h_1, h_2, \ldots, h_H\}$. In the parametrization section, we will allow for the size of rental houses to vary by discount factor type (e.g. h^R_H and h^R_L). This heterogenenity helps account for observed homeownership differences by FD. To finance the purchase of non-rental (owner-occupied) houses, agents borrow using mortgages m'. Importantly, borrowing capacity in the mortgage

[14]Heterogeneity in the discount factor is common in macroeconomics, at least since Krusell and Smith. (1998). However, the modeling and the calibration of β heterogeneity here follows closely Athreya, Mustre-del Río, and Sánchez (2019).

market is endogenously given by a zero-profit condition on lenders due to the limited commitment of agents to repay mortgages.[15]

If agents choose to save in the financial asset $a > 0$, they receive a risk-free rate r. However, when agents borrow ($a < 0$), the discount price of their unsecured debt (q) depends on how much they borrow because debt may be repudiated. Debt repudiation can occur in one of two ways. First, the agent may cease payment. This option is known as delinquency (DQ) or informal default. Importantly, because with delinquency a household's debt is *not necessarily forgiven*, we allow for a probabilistic elimination of debts, with an i.i.d. probability η. This tractably captures not only the absence of a formal elimination of the debt but also the empirical reality that creditors periodically give up on collections efforts.

With probability $1 - \eta$, then, a household's rolled-over debt is not discharged. In this case, the household pays a "penalty" rate, r_R, of interest higher than the average rate paid by borrowers.[16] Moreover, in any period of delinquency, we prohibit saving, and since the agent did not borrow but failed to repay as promised, their consumption equals income. Second, as in standard models of unsecured debt, agents may invoke formal default via a procedure that represents consumer bankruptcy (BK). 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. In what follows, we refer to FD in the model as being in either delinquency or bankruptcy. Since bankruptcy rates are quantitatively much smaller than delinquency rates our results are robust to defining FD as only being in delinquency. However, we think combining both delinquency and bankruptcy is a more holistic measure of distress.

3.2 Nonhomeowners

The options faced by a nonhomeowner with assets a and income y are represented in Figure 5. First, they can choose to either rent or buy a house and become a homebuyer. If renting is chosen, the nonhomeowner must decide between the three options described below. There is a letter associated with each position in the tree, representing the notation

[15] Housing choices, mortgages, and foreclosures are modeled as in Hatchondo, Martinez, and Sánchez (2015).

[16] Athreya, Sánchez, Tam, and Young (2017) analyze facts about informal default and introduced it to heterogeneous-agent models. Athreya, Sánchez, Tam, and Young (2015) use this model to study the effect of the Bankruptcy Abuse Prevention and Consumer Protection Act of 2005.

we use for the value function associated with each choice. For example, the value function for a nonhomeowner with state variables a and y is N. For the sake of brevity, our formal description of this recursive problem is presented in Appendix B.

Figure 5: Decision Tree of a Nonhomeowner

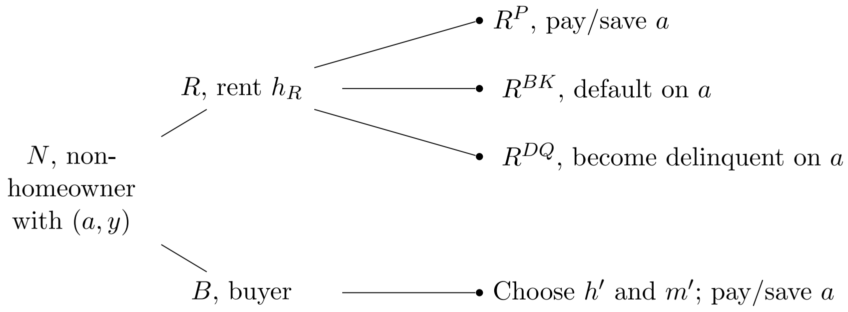

3.2.1 Renting a house

A renter of discount factor type j with income y who decides to *pay* unsecured debt (or has positive financial assets) chooses the next period's financial assets a'. Hence, the agent's budget constraint reads

$$c + q^a_{R,j,n}(a', y)a' = y + a.$$

Here, y denotes income and q^a denotes the price (i.e., discount) applied to financial assets. As noted above, the fact that agents can repudiate debt means that its price will reflect default incentives, which depend on the agent's state vector and amount borrowed a'. For this reason, the function q^a depends on an agent's age/periods left to live (n), type (j), ownership status (in this case renter R), and income y.

Instead, if that renter decides to *formally default* on unsecured debt a, she faces the following trivial budget constraint: $c = y -$ (filing fee), where the "filing fee" is the bankruptcy filing fee.

Finally, if that renter decides to skip payments (i.e., become *delinquent*) on unsecured debt a, they consume $c = y$ and will have financial assets tomorrow equal to

$$a' = \begin{cases} 0, & \text{with prob. } \eta, \\ (1 + r^R)a, & \text{with prob. } 1 - \eta. \end{cases}$$

Here, η is the probability of discharging delinquent debt, and r^R is the roll-over interest rate on delinquent debt.

3.2.2 Buying a house

An agent buying a house must choose the next period's financial assets a', the size of the house h', and the amount to borrow for the house m'. This agent faces the following constraints:

$$\begin{aligned} c + q^a_{j,n}(h', m', a', y)a' &= y + a + q^m_{j,n}(h', m', a', y)m' - I_{m'>0}\xi_M - (1+\xi_B)ph', \\ q^m_{j,n}(h', m', a', y)m' &\leq \lambda ph'. \end{aligned}$$

Here, p is the price of a house, and q^m is the price of a mortgage. The mortgage price depends on the house size, mortgage amount, income, the agent's discount factor type j, and periods left of life n. The second equation is a loan-to-value (LTV) constraint implying that the LTV ratio cannot exceed λ of the value of the house.

3.3 Homeowners

The choices available to an existing homeowner are presented in Figure 6. A homeowner's problem is more complex. On the financial asset dimension, homeowners must decide to default or repay their unsecured debt. On the housing dimension, homeowners can (i) pay their current mortgage; (ii) refinance their mortgage; (iii) default on their mortgage; (iv) sell their house and buy another one; or (v) become a renter. Each option and the associated budget constraint are discussed below.

3.4 Making the mortgage payment

Agents repaying their mortgage who also decide to *pay their unsecured debt* face the following budget constraint:

$$c + q^a_{j,n}(h, m(1-\delta), a', y)a' = y + a - m.$$

Notice that the bond prices these agents face depend on house size h, tomorrow's mortgage size $m(1-\delta)$, the financial assets borrowed or saved a', income, and the agent's discount factor type j. The parameter δ captures the rate at which mortgage payments decay, which may happen, for example, because there is inflation, and payments are fixed in nominal terms.

Agents who pay their mortgage but *formally default on unsecured debt* have the fol-

Figure 6: Decision Tree of a Homeowner

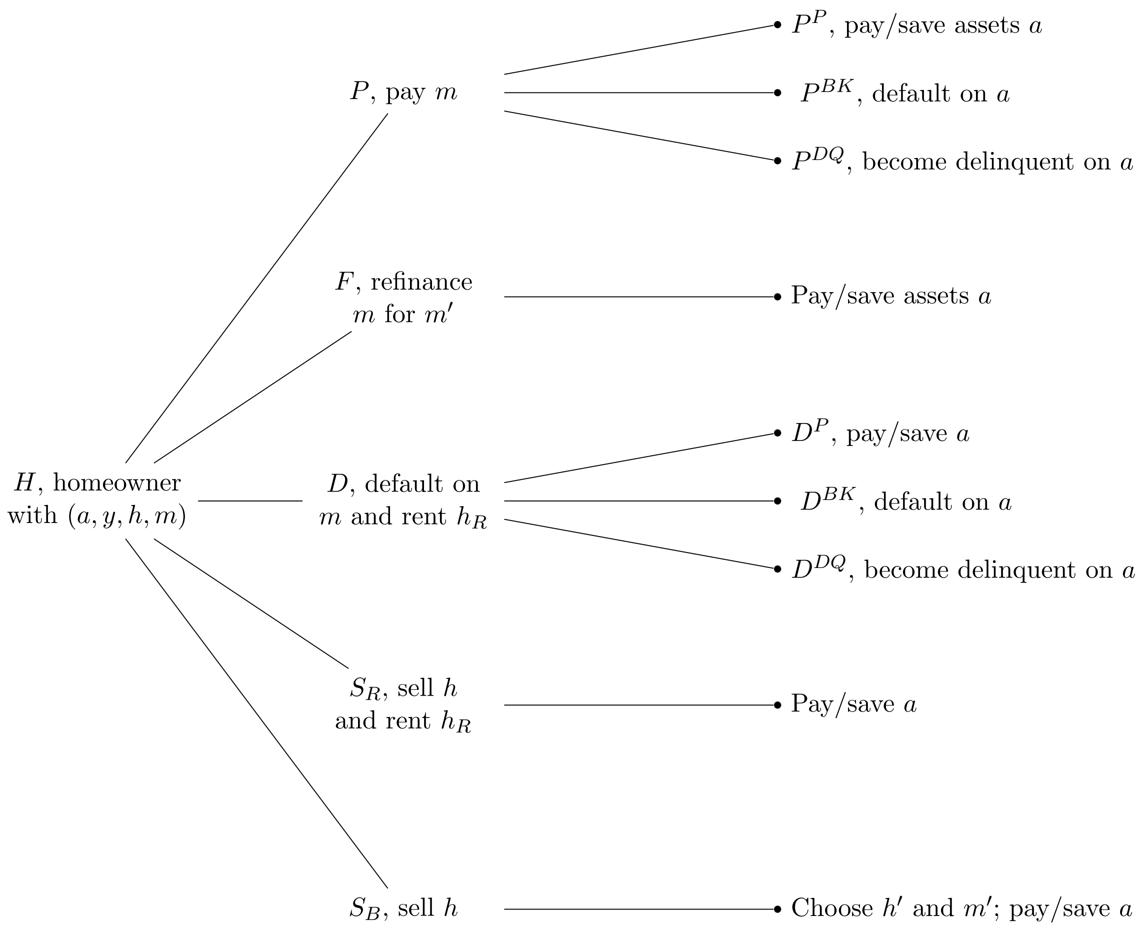

lowing budget constraint, $c = y - (\text{filing fee}) - m$, where "filing fee" is the bankruptcy filing fee and m is the current mortgage payment.

Similarly, households who decide to pay their mortgage but *informally default* on their unsecured debt consume $c = y - m$ and have financial assets tomorrow equal to

$$a' = \begin{cases} 0, & \text{with prob. } \eta, \\ (1 + r^R)a, & \text{with prob. } 1 - \eta. \end{cases}$$

3.4.1 Refinancing the mortgage

An agent who refinances cannot default on unsecured debt a, must prepay their current mortgage, choose next period's financial assets a', and choose the amount to borrow b' with their new mortgage. This problem can be thought of as a special case of a homebuyer who is "rebuying their current home of size h" but who has cash on hand equal to income

y plus financial assets a, minus fees from prepaying their current mortgage m. Thus, the constraints for this problem are:

$$\begin{aligned} c + q^a_{j,n}(h', m', a', y)a' &= y + a - q^*_n m + q^m_{j,n}(h', m', a', y)m' - I_{m'>0}\xi_M, \\ q^m_{j,n}(h', m', a', y)m' &\leq \lambda p h'. \end{aligned}$$

Here, $q^*_n m$ is the value of prepaying a mortgage of size m with n remaining periods worth of payments, which is:

$$q^*_n = \frac{1 - \left(\frac{1-\delta}{1+r}\right)^{n+1}}{1 - \frac{1-\delta}{1+r}}, \text{for} \quad n \geq 1.$$

3.4.2 Foreclosing on the mortgage

An agent who defaults on her mortgage and chooses to *pay her unsecured debt a* immediately becomes a renter and must choose the next period's financial assets a'. Thus, the budget constraint she faces is identical to that of a renter who pays her financial assets: $c + q^a_{R,j,n}(a', y)a' = y + a$.

Using the same reasoning as above, we can write the problem of a mortgage defaulter who chooses *bankruptcy* on unsecured debt as the problem of a renter who files for bankruptcy. Thus, the budget constraint is simply $c = y - filing\ fee$.

Lastly, we can write the problem of a mortgage defaulter who chooses delinquency as the problem of a renter who is also delinquent on existing debt. This means that consumption is given by $c = y$, and financial assets tomorrow are equal to

$$a' = \begin{cases} 0, & \text{with prob. } \eta, \\ (1 + r^R)a, & \text{with prob. } 1 - \eta. \end{cases}$$

3.4.3 Selling the house

A home seller who decides to *rent* cannot default on financial assets. Hence, their optimization problem collapses to that of a renter with financial assets equal to a plus the gains from selling their current house. The agent's budget constraint in this case reads:

$$c + q^a_{R,j,n}(a', y)a' = y + a + ph(1 - \xi_S) - q^*_n m.$$

Here, the term $1 - \xi_S$ is a transaction cost from selling a house with value ph, and $q_n^* m$ is the value of prepaying a mortgage of size m with n periods left.

If instead, the seller decides to *buy another house*, she must also pay her financial obligations. Therefore, this agent's problem is just a special case of a homebuyer with cash on hand equal to income plus current financial assets plus gains from selling the current house. As a result, we can write the constraints for this problem as:

$$\begin{aligned} c + q_{j,n}^a(h', m', a', y)a' &= y + a + ph(1 - \xi_S) - q_n^* m + q_{j,n}^m(h', m', a', y)m' \\ &- I_{m'>0}\xi_M - (1 + \xi_B)ph', \\ q_{j,n}^m(h', m', a', y)m' &\leq \lambda ph'. \end{aligned}$$

3.5 Debt prices

The price of debt, or the interest rate, is determined by risk-neutral lenders that make zero expected discounted profits. In this section, we present the three main components of debt prices. The full specification of each of these (three) prices is in Appendix B.

The *price of a mortgage*, $q_{j,n}^m$, for an agent of type j, with income y, and financial wealth a', for the next period and that promises a payment of m', is given by:

$$q_n^m(h', m', a', y) = \frac{q_{pay,j,n}^m + q_{prepay,j,n}^m + q_{default,j,n}^m}{1 + r},$$

where r is the risk-free interest rate. This equation reveals that the price of a mortgage depends on the likelihood that tomorrow this mortgage will be repaid (first term), prepaid (second term), or defaulted on. Recall, mortgage payment can occur alongside financial debt payment, default, or delinquency. We don't restrict agent choices at all in this regard, which makes our setting very flexible. Meanwhile, mortgage prepayment occurs whenever the agent refinances, sells her current house and rents, or sells her current house and buys another house. In all of these prepayment scenarios, financial debts cannot be repudiated. Lastly, and as is consistent with our overall approach, mortgage default can occur alongside financial debt payment, default, or delinquency. Notice that under this formulation, mortgage prices fully internalize how financial asset positions today and tomorrow affect the probability of mortgage default.

We can express *unsecured* debt prices similarly. When an agent of type j, income y, house size h', and mortgage size m' issues debt and promises to pay a' next period, the amount they borrow is given by $a'q^a_{j,n}(h', m', a', y)$, where:

$$q^a_{j,n}(h', b', a', y) \quad = \quad \frac{q^a_{pay,j,n} + q^a_{DQ,j,n}}{1+r}.$$

First, consider the price of a payment tomorrow, $q^a_{pay,j}$. Payment occurs in a few scenarios. Homebuyers always pay, by assumption. Additionally, payment among homeowners occurs if the owner: (i) is a mortgage payer who also pays her unsecured debt; (ii) is refinancing the mortgage; (iii) is a mortgage defaulter who pays her unsecured debt; (iv) is selling the house to become renter; and (v) is selling the house to buy another house. In these cases, creditors get paid the same amount per unit of debt issued by the household.

Next, consider the price given delinquency tomorrow, $q^a_{DQ,j}$. Among homeowners, this value occurs in two cases: when mortgage payers choose delinquency and when mortgage defaulters choose delinquency. In all of these cases, debt gets rolled over at a rate of $(1 + r^R)$ with probability $(1 - \eta)$. Importantly, though, tomorrow's price of this "rolled-over" debt will depend on the agent's housing status tomorrow. Hence, this reveals that bond prices interact with housing status, as the latter affects the likelihood of financial debt payment, default, and delinquency in the future.

While the bond pricing function of a renter follows a similar form as that of a homeowner, there are some important differences. In particular, a renter has no refinance option nor the option of defaulting on mortgage payments to help make unsecured debt payments. As we'll see in Section 6.3, these differences have important implications for unsecured debt and homeownership.

4 Model estimation and aggregate shock calibration

Before assessing the importance of household heterogeneity in FD and shock dispersion, correlated with FD, in shaping individual and aggregate responses, we take the previously described model to the data in three steps. First, we ensure the model generates the wide dispersion in FD implied by the data. Second, we ask it to generate the observed relationship between FD and aggregate shocks. Lastly, we evaluate the veracity of the model's predictions by comparing its implied MPCs to estimates from the literature.

To accomplish the first two tasks these tasks we take, to our knowledge, a novel approach. We split the US into five different parts. Instead of concentrating on geographical regions (e.g., West, Midwest, Southwest, Southeast, and Northeast), which would have relatively minor differences, we group zip codes in quintiles sorted by the incidence of FD. Thus, zip codes in our groups are not necessarily geographically connected in any way, yet they capture what to us is the critical dimension of similarity: vulnerability to shocks.[17]

We estimate key structural parameters for each of these economies and assign shocks to them consistent with their level of FD. By estimating each economy separately our procedure captures the wide dispersion in FD implied by the data. By assigning shocks to each economy we ensure the entire model (i.e. all five regions) captures the observed positive relationship between FD and aggregate shocks.

4.1 Model estimation

In assigning parameters to each region, we proceed in two steps. First, we directly set values for a subset of the most "standard" parameters and impose that these are common to households across our notion of regions. Second, given these first-stage values, we estimate the remaining parameters so that the model-simulated data matches key statistics on wealth, home ownership, and FD for each of the five economies.

4.1.1 Assigning first-stage parameters

Table 1 collects the parameters set externally. A period in the model refers to a year.[18] Agents enter the model at age 25, retire at age 65, and die no later than age 82. We set the risk-free interest rate at 3 percent. In addition, we externally calibrate the parameters governing the income process, bankruptcy filing costs, retirement, and mortality. As

[17] Of course, precisely due to the effective selection into economically similar groups, our chosen partition of the data precludes general equilibrium analysis *inside* each group. That is, the spillovers across groups would be very significant. Nonetheless, to alleviate the concern of spillovers across zip codes, we redid our exercises by grouping counties instead of zip codes, and the results are similar.

[18] In contrast, in the data we measure delinquency as being 30+ days overdue. While in principle there is a discrepancy between data and model, in practice, this difference shouldn't matter much as delinquency is fairly persistent. In our Equifax sample, conditional on being delinquent in one quarter during a year, individuals are likely to be delinquent at some other point in the year. Intuitively, being delinquent in one month restricts credit access henceforth. Our model accounts for this by precluding any borrowing during the entire year when in delinquency. Additionally, whether we measure delinquency at 30+ or 120+ days overdue, the conditional probabilities of being delinquent in the future are very similar at all time horizons (e.g. 2 years later, 4 years later, etc.). Thus, the persistence of delinquency or FD (which is key for parameter identification) is very similar regardless of the time period definition.

discussed below some preference and housing parameters are also externally set. The initial distribution of net financial wealth-to-earnings is set to match the distribution of net financial wealth-to-earnings of 25-year-olds in the Survey of Consumer Finances between 1998 and 2016.

For time preference, we follow Athreya, Mustre-del Río, and Sánchez (2019) in assuming that agents can either discount the future relatively little (i.e., be "patient") and have discount factor β_H, or discount it more significantly (i.e., be "impatient") and use discount factor β_L. This heterogeneity allows the model to capture well the joint distribution of net financial wealth, delinquency (incidence and persistence), and bankruptcy. We set β_H=1.00 and β_L=0.80, which is within the range of estimates in Athreya, Mustre-del Río, and Sánchez (2019) and also Aguiar, Bils, and Boar (2020).

In terms of preferences for consumption and housing, we assume households experience utility with a constant elasticity of substitution:

$$u(c,h) = \frac{((1-\theta)c^{1-1/\alpha} + \theta h^{1-1/\alpha})^{(1-\gamma)/(1-1/\alpha)}}{1-\gamma},$$

where γ denotes the risk aversion parameter, α governs the degree of intra-temporal substitutability between housing and nondurable consumption goods, and θ determines the expenditure share for housing. Following Hatchondo, Martinez, and Sánchez (2015), we set γ to 2, α to 0.5. The value of α is consistent with estimates from Hanushek and Quigley (1980), Siegel (2008), and Li, Liu, Yang, and Yao (2015). We set θ to 0.14, which matches the share of housing in total consumption expenditures in NIPA data.[19]

Since our model must match as well as possible the overall homeownership rate and the joint distribution of homeownership and FD, we assume that the size of rental houses h^R differs by preference type. The size of rental houses for L-types is denoted as h_L^R and the size of rental houses for H-types as h_H^R. Differences in these two parameters help capture differences in the utility of homeownership (or disutility of renting) across types in a succinct fashion. Given the combinations of homeownership rates and incidence of FD that the data display, our model implicitly requires a very high degree of homeownership (near 100 percent) among patient types across all quintiles of FD. Thus, we set h_H^R to a very low value and leave h_L^R as a parameter to be determined below.

[19] A similar calibration strategy is done in Jeske, Krueger, and Mitman (2013).

Turning to owner-occupied houses we set a few more parameters using external information. Because median home value to income ratios do not vary dramatically across quintiles of FD, we set house prices constant across the five economies at $p = 3.3$.[20] Given the sizes of houses for purchase, this value helps generate median home value-to-income ratios between 3.2 and 3.3 as observed in the data. Next, we assume the mortgage payments decay rate is δ=0.02, so that mortgage payments decay with the average rate of inflation. Finally, we allow for owner-occupied houses to be subject to appreciation and depreciation shocks. The depreciation shocks are such that, on average, the owner-occupied housing stock depreciates at 1.5 percent annually, following Kaplan, Mitman, and Violante (2020). Conditional on not depreciating, a house can appreciate. These shocks are such that, on average, home values increase by 3 percent annually. This roughly matches the average annual growth rate of house prices (relative to CPI) based on the Case-Shiller National House Price Index.

Next, following Livshits, MacGee, and Tertilt (2007), the penalty rate for delinquent debt is set at 20% annually, and the bankruptcy filing costs are at 2.8 percent of average income, or roughly \$1,000. We set the discharge rate of delinquent debt to 25 percent annually, so $\eta = 0.25$. This is within the range of estimates reported in Athreya, Mustre-del Río, and Sánchez (2019).

Turning to the income-process parameters, we consider restricted income profile (RIP)-type income processes following Kaplan and Violante (2010). During working ages, income has a life-cycle component, a persistent component, and an i.i.d. component:

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

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

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

We assume $\epsilon^i_{n,t}$ and $e^i_{n,t}$ are normally distributed with variances σ^2_ϵ and σ^2_e, respectively. While the income process does not vary across quintiles of FD, the level of income does.[21]

[20] Note, we assume rental houses are free to ensure everyone in our model can afford some housing. A similar assumption is made in Hatchondo, Martinez, and Sánchez (2015).

[21] Alternatively, one could hypothesize the key difference across quintiles of FD is heterogeneity in the income process. We estimated an alternative version of our baseline model where the variance of transitory (σ^2_ϵ) and persistent (σ^2_e) shocks to income varied by quintile of FD, but all other parameters were constant across quintiles and all other sources of heterogeneity were shut-down (e.g. $s_L = 0$,

We normalize the level of income across quintiles such that the level of income in the third quintile of FD is equal to 1. These normalizations imply that income in the first quintile of FD is about 40 percent larger than the third quintile. Meanwhile, income in the fifth quintile of FD is about 24 percent smaller than in the third quintile.

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: $\max\{A_0 + A_1 exp(z^i_{W1}), A_2\}$. 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 percent, with an average replacement rate of 47 percent. The age-specific survival probabilities follow Kaplan and Violante (2010).

Table 1: Externally Set Parameters

Parameter	Value	Definition	Basis
l	—	Life-cycle component of income	Kaplan and Violante (2010)
W	65	Retirement age	U.S. Social Security
ρ_n	—	Mortality age profile	Kaplan and Violante (2010)
a_0	—	Initial net financial asset distribution	SCF 1998-2016
σ^2_ϵ	0.05	Variance of ϵ	Kaplan and Violante (2010)
σ^2_e	0.01	Variance of e	Kaplan and Violante (2010)
r	0.03	Risk-free rate	Standard
γ	2	Risk aversion	Standard
α	0.5	Elasticity of substitution	Standard
θ	0.11	Consumption weight of housing	Hatchondo et al. (2015)
ξ_B	0.03	Cost of buying a house, households	Gruber and Martin (2003)
ξ_S	0.03	Cost of selling a house, households	Gruber and Martin (2003)
$\bar{\xi}_S$	0.22	Cost of selling a house, banks	Pennington-Cross (2006)
ξ_M	0.15	Cost of signing a mortgage	U.S. Federal Reserve
δ	0.02	Mortgage payment decay	Average inflation
A_0	0.7156	Replacement ratio	U.S. Social Security
A_1	0.04	Replacement ratio	U.S. Social Security
A_2	0.14	Replacement ratio	U.S. Social Security
λ	0.9	LTV limit	Positive down payment
f	0.028	Cost of filing for bankruptcy/mean(inc)	Livshits et al. (2007)
r_R	0.2	Roll-over rate on delinquent debt	Livshits et al. (2007)
β_H	1.00	Discount factor of patient types	Athreya et al. (2019)
β_L	0.80	Discount factor of impatient types	Athreya et al. (2019)
h^R_H	0.001	Size of rental house for patient types	See text
p	3.33	House prices	See text
η	0.25	Discharge rate of delinquent debt	See text

$\beta_L = \beta_H = \beta$, $h^R_L = h^R_H = h^R$). Under this estimation strategy, the model performed considerably worse in matching the empirical targets. Additionally, the estimation implied implausibly large (relative to existing estimates in the literature) parameter values for the variance of both income shocks.

4.1.2 Estimating the remaining parameters

The only remaining parameters to be determined are (i) the share of impatient types in the population s_L; and (ii) the rental house size h_L^R for impatient types. We estimate these two parameters so that the model replicates some critical features of the data on homeownership, financial wealth, and FD for each of the five regions we construct.

Table 2 presents the model's fit for each of the quintile-specific moments. The model does a good job of matching differences in financial wealth across the five quintiles, though it cannot quite reproduce the extreme differences between Q1 and Q5. Additionally, it replicates the fact that homeownership declines as FD rises and matches the share of individuals in FD that have housing debt well. Because most individuals in FD who own a home will tend to have mortgages or home equity lines of credit (HELOCs), this measure can be thought of as a good proxy for the homeownership rate conditional on being in FD.

The rest of the table focuses on FD (recall defined as being in either delinquency or bankruptcy) and shows that the model does well at reproducing the overall patterns. The model matches closely the fact that average delinquency rates rise with each quintile of FD. The model also matches the fact that bankruptcy rates rise with FD, but not as steeply. For example, in the data, the average bankruptcy rate in Q5 is roughly 1.6 times that of Q1. Meanwhile, the corresponding calculation in the model is 1.5.

Turning to the persistence of FD the model also matches that it tends to fall over time within a given quintile and, perhaps counterintuitively, also tends to fall across quintiles as FD increases. Here the persistence measure is defined as a relative ratio. For example, the persistence measure over two years is defined as the ratio between (i) the probability of being FD in two years conditional on being in FD today, and (ii) the probability of being FD in two years conditional on not being in FD today. Overall, the model matches the fact that the persistence of FD gradually falls with the time horizon.

Table 2: Model Fit by Quintile of FD

	Q1		Q2		Q3		Q4		Q5	
	Data	Model	Data	Model	Data	Model	Data	Model	Data	Model
Savings/Inc	2.44	1.73	1.96	1.53	1.78	1.35	1.57	1.23	1.06	1.03
Home ownership*	76.3	76.63	71.93	71.04	68.76	62.34	64.25	61.63	61.69	52.58
Housing leverage	44.11	30.06	47.98	36.56	44.57	40.89	46.04	44.34	43.36	44.74
Housing debt> 0*	49.77	28.75	44.67	28.72	39.83	25.04	36.27	27.43	31.84	24.2
Housing debt> 0* conditional on FD	33.31	32.18	30.72	30.75	28.37	20.89	26.9	27.44	25.99	21.37
Housing debt/Inc	1.47	0.86	1.57	1.03	1.57	1.17	1.59	1.31	1.48	1.52
Mortg def rate*	1.52	1.3	1.81	1.69	2.24	2.2	2.58	2.18	3.34	2.49
DQ rate*	8.98	10.22	12.65	13.4	15.43	16.22	18.28	18.55	23.93	22.3
BK rate*	0.39	0.59	0.55	0.67	0.63	0.65	0.65	0.69	0.64	0.66
Persistence of FD:										
Over 2 yrs	9.2	6.5	8.05	5.31	6.82	4.87	5.89	4.03	4.83	3.58
Over 4 yrs	6.15	5.3	5.36	4.21	4.57	3.74	3.99	3.17	3.2	2.77
Over 5 yrs	5.36	5.31	4.63	4.18	3.98	3.61	3.45	3.09	2.86	2.68
Over 6 yrs	4.86	5.29	4.17	4.15	3.57	3.58	3.2	3.07	2.58	2.64
Over 8 yrs	3.89	5.13	3.56	4.11	2.95	3.52	2.61	3.01	2.19	2.59
Over 10 yrs	3.4	4.38	3	3.68	2.66	3.23	2.37	2.86	2.05	2.47
SSE	1.15		0.66		0.56		0.39		0.35	

Notes: $*$ in percent. SSE is the sum of squared errors for each quintile. "Savings/Income" represents mean net financial wealth divided by mean income, and "With housing debt / In FD" is the percent of the population with housing debt, conditional on being in FD.

Table 3 shows the resulting parameter estimates and reveals significant and systematic differences across quintiles of FD. Most notably, the share of impatient individuals rises from Q1 (least distressed) to Q5 (most distressed). In Q1, 30 percent of the population is impatient and discounts the future relatively more. In Q5, by contrast, nearly 60 percent of the population is impatient. Thus, between Q1 and Q5, there is nearly a doubling of this share. The model requires this difference in the type-L population between Q1 and Q5 to match similarly significant differences between these quintiles in the data. Both delinquency and bankruptcy rates are higher in Q5 than in Q1. In contrast, homeownership is nearly 15 percentage points lower in Q5 versus Q1. Lastly, net financial wealth to income is less than half as big in Q5 compared to Q1. A more significant share of impatient types in Q5 helps to generate these features.

The model estimates also imply significant heterogeneity in rental house sizes within and across quintiles. Focusing first on the within differences, recall that for all quintiles, the size of rental houses for type-H individuals is close to zero, by assumption. The estimates in the middle of Table 3 therefore allow us to quickly reject the null of no differences in rental house size between types, regardless of quintile of FD. Across quintiles, the estimates also allow us to reject the null of equal rental house sizes for type-L individuals who live in Q1 versus Q5. Interestingly, the model requires a smaller value of h_L^R in Q5 versus Q1.

What helps identify this parameter from the data? Within-quintile differences in h_L^R versus h_H^R depend on the within-quintile difference in the share of individuals with housing debt conditional on being in FD and the overall homeownership rate. Treating the former as a proxy for homeownership among those in FD suggests that individuals in FD are less likely to own homes than the average person. Even though individuals in FD are less likely to own a home (compared to those not in FD) because of their prior financial choices, the model still requires differences in h^R to match the extreme differences in ownership between those in FD and not. Taking Q1 as an example, individuals in FD are roughly 43 percentage points less likely to own a home than the average individual in this quintile.

Cross-quintile differences in h_L^R crucially depend on cross-quintile differences in the aforementioned measure (i.e. the difference between the share of individuals with housing debt conditional on being in FD and overall homeownership). As can be seen in Table 2, this measure ranges from -43 percentage points in Q1 to -36 percentage points in Q5. Thus, from the model's perspective, conditional on FD, homeownership is relatively less appealing in Q1 than in Q5. To replicate this pattern the model requires a higher value for h_L^R in Q1 than in Q5.

A more general question is which heterogeneity (in discount factors or rental house sizes) is most important to generate our results. In Appendix C, we present results for two restricted models. In one we shut down differences in rental house sizes within and across

Table 3: Parameter Estimates by Quintile of FD

Parameter	Q1	Q2	Q3	Q4	Q5
s_L	0.29	0.37	0.45	0.50	0.58
	(0.09)	(0.05)	(0.04)	(0.05)	(0.03)
h_L^R	4.61	3.77	3.92	2.99	2.95
	(0.07)	(0.02)	(0.06)	(0.01)	(0.01)

Notes: Asymptotic standard errors appear in parentheses.

quintiles (i.e. $h_L^R = h_H^R = h^R$ across quintiles). In the other, we shut down differences in discount factors, also within and across quintiles (i.e. $\beta_L = \beta_H = \beta$ across quintiles). In each case, the fixed parameter is set to the quintile average obtained from Table 3.

Those results reveal that each type of heterogeneity plays a different role in matching the empirical targets. Discount factor heterogeneity is necessary to match the dispersion in savings/income ratios, the level and persistence of delinquency rates, and housing debt/income ratios we observe across quintiles of FD. Without this heterogeneity savings/income ratios and delinquency rates (levels and persistence) are too low, while house debt/income ratios are too high. Heterogeneity in rental house sizes helps match the dispersion in bankruptcy rates and homeownership rates we observe across quintiles of FD. Without this heterogeneity homeownership and bankruptcy rates are too high.

Naturally, neither restricted model performs as well as our baseline model. In particular, neither restricted model matches the percentage of individuals with housing debt conditional on being in FD. Critically, allowing for both types of heterogeneity (discount factor and rental size) and for L-types to be both impatient and have larger rental sizes allows the model to quantitatively match the fact that individuals in FD are less likely to own homes. While both models predict that individuals in FD are less likely to own homes than the average individual, the predicted magnitudes are substantially off. In the data, the average "homeownership gap" between individuals in FD versus the average individual is -40 percentage points. The model with only discount factor heterogeneity implies a gap of -11 percentage points, while the model with only rental house size heterogeneity implies a gap of only -3 percentage points. In contrast, the baseline model implies a gap of -38 percentage points, which is very close to the data.

However, despite these large differences in empirical fit, many of the empirical patterns are reasonably approximated with discount factor heterogeneity alone. Additionally, we show in Section 5 that model-implied MPCs, and their relationships with FD, are fairly similar between the baseline model and the one restricted to only allow for discount factor heterogeneity.

4.2 Aggregate shock calibration

Having estimated five different economies to capture the wide dispersion of FD across the US, we now focus on replicating the correlation between FD and aggregate shocks observed in the GR and CV19 pandemic. To do so, we create two stylized recessions that mimic how shocks were distributed across FD regions. The first is an unexpected permanent decline in house prices, similar to that observed during the GR. Since houses are assets and estimates of an autoregressive process for prices are very close to a random walk, we assume house price shocks are permanent. The second is an unexpected temporary decline in labor income, similar to the CV19 pandemic. Since most of the effect of the pandemic on labor earnings in the US was short-lived, we assume that these income shocks are temporary. In both cases, our quantitative analysis treats these shocks as exogenous and is not meant to capture all the features of these downturns. Instead, our goal is to understand how aggregate shocks transmit into consumption changes when FD is an option and when it is correlated with shock exposure.[22]

Table 4 shows the shocks buffeting each quintile of FD and reveals, by construction, significantly different experiences across quintiles for each of the considered downturns. In terms of house-price shocks, we use the data presented in Section 2.2 to calculate the average change in house prices between 2007 and 2008 for each quintile.[23] In terms of labor earnings shocks, we construct the distribution of earnings losses from survey evidence from Bick and Blandin (2021).

These distributions highlight the positive relationship between aggregate shocks and FD. In terms of house price declines, the Q1 economy (lowest FD) only experienced a 7% decline in house prices. Meanwhile, the Q5 economy (highest FD) experiences an 11.5% decline in house prices. For labor earnings declines, the disparity is even starker. Focusing

[22]For a rich analysis of the decline in house prices observed during the GR see Garriga and Hedlund (2017).

[23]We obtain very similar results using the average yearly change between 2006 and 2009 as well.

Table 4: Calibration of House-price and Labor-income Shocks

FD Quintile	Average decline in house prices	Percent of population with earnings loss of: 0%	25%	50%	Average earnings loss
1	7.0	80.3	5.3	14.4	8.5
2	8.6	79.3	5.6	15.1	9.0
3	10.0	78.2	5.1	16.7	9.6
4	10.9	76.5	5.9	17.6	10.3
5	11.5	72.4	5.9	21.7	12.3

Sources: Zillow and Bick and Blandin (2021).

on severe earnings losses (a 50% decline relative to pre-shock earnings), roughly 14% of households in Q1 received this type of shock, compared to nearly 22% of households in Q5.

5 Model validation via MPCs

Before exploring the quantitative implications of our model we evaluate the model's performance in replicating external information on consumption responses and provide some additional analysis on the model's mechanics. We show the model's aggregate MPCs line up with external estimates. Additionally, we confirm the model generates a systematic relationship between FD and MPCs out of house-price shocks that also aligns with the empirical evidence shown earlier in Section 2.1. We show that most of this prediction is obtained in a model with only discount factor heterogeneity, suggesting this is the key dimension of ex-ante heterogeneity in the model. Finally, we explore the usefulness of looking at FD when measuring MPCs. Model-based exercises suggest that much of the dispersion in MPCs by ex-ante heterogeneity type (which is unobservable) can be similarly identified by differences in FD across individuals (which is easily observable). In other words, categorizing individuals by FD is an empirically tractable way of categorizing individuals by MPC and unobservable type.

5.1 Model versus data

We now verify the degree of transmission of shocks (to either house prices or income) into consumption in each of the five economies is consistent with external evidence. To do so, we present model-implied MPCs out of house-price and labor-income shocks. In all cases, the MPCs are the response of non-housing consumption to the corresponding shock. The similarities with empirical estimates that we find are reassuring, providing empirical support to the quantitative claims we make in the next section.

First, we consider how consumption responds to the house-price shocks described in the previous section. It takes time for consumption in each quintile to adjust after a permanent shock. To capture the change over time, we calculate the average annual MPC over three years following the housing shock.[24] The results are displayed in the first two rows of Table 5.

Table 5: Model-implied MPCs

	Aggregate	Q1	Q2	Q3	Q4	Q5
House-price shocks, all	0.071	0.065	0.074	0.067	0.076	0.078
House-price shocks, owners	0.088	0.083	0.086	0.086	0.095	0.095
Labor-income shocks	0.299	0.242	0.274	0.314	0.334	0.384

Notes: This table presents marginal propensities to consume (in dollars) out of permanent house-price shocks and transitory income shocks. In the case of labor-income shocks, the calculation is restricted to working-age individuals. The quintiles are based on the distribution of FD, with Q1 being the least distressed quintile and Q5 the most.

The first cell of this table shows an aggregate MPC of 7 cents per dollar. This number is virtually identical to the IV estimate of 7.2 cents per dollar that Mian, Rao, and Sufi (2013) obtain in their county-level analysis of the MPC to consume out of housing wealth shocks. This headline number includes both homeowners and non-homeowners. If we restrict our attention to homeowners the model-implied aggregate MPC rises to 8.8 cents per dollar, as seen in the second row, first column. The difference between these two MPCs suggests that homeowners drive the bulk of the MPC and that renters have a negative MPC in response to house-price shocks. That homeowners drive most of the aggregate change in consumption aligns with Aladangady (2017). The negative MPC among renters comes from the fact that in the model they will eventually become homeowners. Thus, when house prices unexpectedly decline, they experience a small positive income effect that allows them to consume more while still purchasing houses as planned. In contrast, established homeowners experience a negative wealth effect and thus decrease their consumption.

The others columns of the top two rows highlight the role of FD in shaping the model-

[24]Calculating MPCs for shorter or longer time horizons does not alter our conclusions.

implied MPCs out of housing. The top row shows that when looking at all households, the MPCs out of housing shocks range from 6.5 cents per dollar (Q1) up to 7.8 cents per dollar (Q5). However, because homeownership rates systematically fall with FD and since we have already noted that the model implies negative MPCs among renters, the cross-quintile comparison of the first row masks some more salient trends.

Restricting our attention to homeowners, as in the second row, reveals more systematic differences in MPCs by quintile of FD. We still observe dispersion in MPCs across quintiles of FD ranging from 8.3 cents per dollar (Q1) to 9.5 cents per dollar (Q5). Importantly, however, the model delivers MPCs out of housing shocks that rise with FD, consistent with the evidence presented in Section 2.1. In a similar vein to MPCs rising with FD, Aladangady (2017) finds that MPCs rise with debt-service ratios (DSRs), from essentially zero (among households with below median DSRs) to 0.127 (among households with above median DSRs).

The bottom row of Table 5 shows that the model also performs well in generating realistic MPCs out of labor-income shocks. Here we focus on working-age individuals, as those in retirement do not receive the shock. The model implies an aggregate MPC of \$0.30 per \$1 transitory increase in income. This MPC is similar to that in Sahm, Shapiro, and Slemrod (2010), who report "an aggregate MPC after one year of about one-third." The size of this response is also close to empirical estimates like those in Coronado, Lupton, and Sheiner (2005) and Jappelli and Pistaferri (2006).

The results for individual quintiles show there is significant heterogeneity behind this aggregate number, where higher FD is associated with a higher MPC out of labor-income shocks. In particular, the difference in MPCs between the least and most distressed quintiles (24 vs. 38 cents) is in line with the empirical evidence presented in Parker (2017) that households with low liquidity spend at a significantly higher rate than that of high liquidity households.

5.2 Which heterogeneity matters

Table 6 explores how the specification of heterogeneity shapes the model-implied MPCs. Each panel of this table displays MPCs for the baseline model and the two previously mentioned restricted models: a version of the baseline model where only discount factor heterogeneity is allowed (β-het model); and, a version of the baseline model where only heterogeneity in rental house sizes is allowed (h^R-het model). While the aggregate MPCs are very similar regardless of the specification of heterogeneity, the distribution of MPCs by quintile of FD does depend on the heterogeneity specification. Overall, the analysis suggests discount factor heterogeneity is crucial.

The top panel shows that aggregate MPCs out of house-price shocks are similar across models, but this masks differences across the distribution of FD. The first column of this panel shows both restricted models deliver MPCs that are within 0.5 cents of the baseline MPC. However, the remaining cells of this panel show the models imply different relationships between FD and the MPC out of housing shocks. The baseline model suggests MPCs weakly increase with FD, ranging from 6.5 to 7.8 cents (between Q1 and Q5). The β-het model delivers a slightly steeper relationship with FD, with MPCs ranging from 6.4 to 8.5 cents. The h^R-het model, in contrast, suggests MPCs are slightly decreasing with FD, with MPCs ranging from 8 to 4.7 cents.

These conclusions are strengthened when we restrict our attention to homeowners, as is done in the second panel of Table 6. Among owners, the overall MPCs are still fairly similar across models, as seen in the first column of this panel. The second row of this panel shows that in the β-het model, MPCs among owners increase sharply with FD, ranging from 5.9 to 10.7 cents. Again and in contrast, in the h^R-het model, MPCs among owners decrease with FD, ranging from 9 to 5.1 cents. Overall, this suggests that discount factor heterogeneity is key in generating the pattern of MPCs out of house-price shocks increasing with FD that the data suggests and baseline model replicates.

The last panel of Table 6 shows that aggregate MPCs out of labor-income shocks are also similar across models, but this again masks differences by quintile of FD. The first column of this panel shows that the restricted models deliver MPCs within 3 cents of the baseline MPC of 30 cents. The remaining cells show that only the β-het model implies an increasing relationship between MPCs and FD like what the baseline model predicts. Indeed, the second row of this table shows that the β-het model delivers MPCs that range from 24.6 to 35.7 cents. In contrast, the third row shows that the h^R-het model implies a much flatter relationship between FD and MPCs, with MPCs ranging from 26.6 to 27.6 cents. Again, this suggests that discount factor heterogeneity is what generates MPCs out of labor-income shocks increasing with FD. This model-based observation is consistent with the empirical work of Parker (2017) that argues that "the majority of lack of consumption smoothing is predicted by a simple measure that can be interpreted as impatience."

Table 6: Model-implied MPCs by Heterogeneity Specification

	Aggregate	Q1	Q2	Q3	Q4	Q5
House-price shocks, all						
Baseline model	0.071	0.065	0.074	0.067	0.076	0.078
β-het model	0.073	0.064	0.071	0.074	0.077	0.085
h^R-het model	0.076	0.080	0.074	0.079	0.092	0.047
House-price shocks, owners						
Baseline model	0.088	0.083	0.086	0.086	0.095	0.095
β-het model	0.078	0.059	0.068	0.078	0.092	0.107
h^R-het model	0.084	0.090	0.077	0.091	0.102	0.051
Labor-income shocks						
Baseline model	0.299	0.242	0.274	0.314	0.334	0.384
β-het model	0.289	0.246	0.271	0.297	0.314	0.357
h^R-het model	0.270	0.266	0.271	0.265	0.276	0.276

Notes: See Table 5. The only β-het rows present the results of a restricted version of the baseline model where individuals only differ in their discount factors and are offered rental houses of the same size. The only h^R-het rows present the results of a restricted version of the baseline model where individuals only differ in the size of rental houses offered, but their discount factors are the same.

5.3 The importance of FD

So far, we have only explored how MPCs differ across quintiles of FD. However, we have yet to assess whether MPCs differ by FD in the expected way. Additionally, we have yet to uncover whether differences in MPCs by FD, which is observable, are similar to differences in MPCs by individual type, which is unobservable. Any similarity between these two sets of MPCs would provide further evidence that FD is a useful metric since it ultimately reflects ex-ante unobservable differences across individuals that drive differences in MPCs.

Table 7 confirms that MPCs are larger for individuals in FD (regardless of shock). To circumvent differences in MPCs arising from the life cycle profile of homeownership, the analysis in this table focuses on individuals ages 30 to 50. As can be seen by the first column of the table, for this age group the housing MPCs are somewhat smaller than for the overall population, whereas the income MPC is slightly larger. The next two columns of this table show that these average MPCs mask substantial differences by prior FD status. The second column of this table presents MPCs for individuals who in steady state (i.e. absent the the shock in question) are in FD, while the third column presents MPCs for individuals who in steady state are not in FD. The second row shows that owners in prior FD have an MPC out of permanent housing shocks of 11 cents, which is more than twice the MPC among owners not in prior FD. Turning to labor-income shocks, the third row reveals that individuals in prior FD have an MPC of 69 cents, which is also more than twice as large as the MPC among individuals not in prior FD.

The differences in MPCs by FD are large, but, as the last two columns of Table 7 show, they are consistent with the large differences in MPCs by ex-ante type. Categorizing by individual type suggests that type-H owners (of this age group) essentially don't respond to house-price shocks. Rather, type-L owners respond quite strongly, reducing their consumption by roughly 19 cents per dollar of home value lost. Turning to labor-income shocks, type-H workers also respond very modestly to these shocks. In contrast, the MPC for type-L workers is 67 cents, which is very similar to the 69-cent MPC among workers in prior FD. In sum, these simple calculations suggest not only is FD a good way to capture meaningful differences in MPCs across individuals, but these differences appear to reflect underlying (and unobservable) differences in preferences.

Table 7: Model-implied MPCs by FD and heterogeneity type

	all	in FD	no FD	Type-L	Type-H
House-price shocks, all	0.034	-0.013	0.041	0.139	-0.008
House-price shocks, owners	0.057	0.111	0.047	0.189	-0.007
Labor-income shocks	0.331	0.689	0.268	0.672	0.090

Notes: This table presents marginal propensities to consume (in dollars) out of permanent house-price shocks and transitory income shocks for individuals ages 30 to 50. The columns labeled “in FD” and ”no FD” refer to individuals who in the steady state (i.e. absent the corresponding shock) are in FD or not, respectively.

6 Quantitative results

In this section, we assess the importance of FD for various outcomes. First, we examine how FD shapes the aggregate and cross-sectional responses of consumption to shocks. Second, we consider how the inclusion of FD interacts with the housing market, paying special attention to foreclosure and housing leverage.

We find that FD matters for aggregate and cross-sectional consumption since it *alters* (i.e. either amplifies or attenuates) the response of the aforementioned measures by nearly 25 percent, on average. FD modifies responses to house-price shocks by more (around 30 percent), and responses to labor-income shocks by less (around 20 percent). Much of this role of FD is not a mechanical consequence of the correlation between prior FD and aggregate shocks. Rather, the attenuation or amplification of FD identified in these exercises has to do with FD as a form of debt repudiation.

We also find that FD has important implications for the housing market. Disallowing FD (described below) leads to lower foreclosure rates, housing leverage, housing debt-to-income ratios, and homeownership. These results are largely related to how FD as an another margin of debt adjustment alters the composition of homeowners. In particular, allowing for FD is associated with higher homeownership rates among impatient types and their presence in the pool of homeowners boosts the overall mortgage default rate.

6.1 The role of FD

Gauging the amplification (or attenuation) effect of FD can only be done relative to a model without FD. Here we consider an alternative model with no FD and, therefore, no differences in shocks correlated with FD. This alternative model is a heterogeneous agent model with incomplete markets, housing choice, and mortgages. Relative to our baseline model it precludes any notion of FD by imposing an ad hoc borrowing constraint, which is assumed to be the same for all agents. With no notion of FD, each shock in this model matches the average decline documented in Table 4. To make comparisons as clear as possible, we keep the same distribution of preferences in this alternative model as in the baseline model. Our main quantitative conclusions are essentially unchanged if we re-estimate the alternative model to match all the non-FD related moments described in Section 4.1.2. More details on each of these alternative models are provided in Appendix

D.

Table 8 shows the importance of FD in shaping the responses of aggregate and individual consumption to shocks. For the aggregate, we examine the change in consumption in response to each shock. For the cross-section, we consider two measures. The first is the change in the p90/p10 ratio of the consumption distribution. This gives us a sense of how the response of consumption inequality to shocks is shaped by FD. The second measure is how consumption-based poverty changes. This helps us assess how FD shapes the response of consumption among those in very precarious spending states. The first column of the table collects all of the measures of consumption responsiveness we consider under the baseline model. The second column displays the corresponding numbers implied by the alternative economy. The third column presents the difference in responsiveness between the two models, or the amplification effect of FD, as a percentage of the baseline model response. The fourth column presents the resulting amplification effect if the alternative model is re-estimated. The top panel focuses on house-price shocks, while the bottom panel pertains to labor-income shocks.

Table 8: The amplification role of FD by shock and consumption measure

	Baseline model (1)	Alternative model (2)	Amplification (3)	Amplification (re-estimated) (4)
House-price shocks				
Consumption p90/p10	-4.85	-3.92	19.25	50.31
Consumption-based poverty	-3.06	-1.62	47.09	44.56
Aggregate consumption	-1.80	-1.36	24.50	10.35
Average absolute value			30.28	35.07
Labor-income shocks				
Consumption p90/p10	14.63	16.22	-10.89	-5.34
Consumption-based poverty	16.39	21.33	-30.17	-45.61
Aggregate consumption	-3.31	-2.78	15.85	8.26
Average absolute value			18.97	19.73

Notes: Columns (1) and (2) are measured as percent deviations from steady-state. In the housing shock case, these are average changes over three periods following the shock. In the labor-income shock case, the change is measured only in the period of the shock and is calculated over the working-age population. Columns (3) and (4) measure amplification as a percentage of the corresponding value in Column (1). In Column (3), the amplification is calculated between the baseline model (1) and the alternative model (2). In Column (4), the amplification is calculated between the baseline model (1) and the re-estimated alternative model (unreported). Details of the alternative and re-estimated alternative models appear in Appendix D. The average absolute value represents the average of the absolute values of either Column (3) or (4).

The first row shows that FD amplifies the effect of house prices on consumption inequality. Following a fall in house prices, consumption inequality *declines* in both the baseline and alternative models, but the drop is larger in the former. Hence, allowing for FD amplifies the drop in consumption inequality by nearly 20 percent, as shown in Column (3). Much of the reduction in inequality (in either model) comes from the 10th percentile of consumption increasing; i.e. consumption increasing at the bottom of the distribution.

Consistently, the second row of this table shows that a fall in house prices also reduces consumption-based poverty and the fall is larger in the baseline model. We follow Meyer and Sullivan (2019) by targeting a consumption-based poverty threshold of 13 percent in steady state (the average poverty rate in the U.S. between 2015-2018) and measure how the population share below this threshold changes in response to each shock.[25] In the baseline model, poverty declines roughly 3 percent (i.e. roughly 40 basis points relative to the steady state threshold of 13 percent), which is nearly double the drop in poverty implied by the alternative model without FD.

The greater drop in inequality and poverty in the baseline versus alternative model is largely explained by the differing composition of individuals at the left tail of the consumption distribution in each model's steady state. In both models, individuals at the left tail are more likely to be renters, who potentially benefit from the cheaper houses. However, among these renters, the baseline model has a larger share of type-H individuals compared to the alternative model (28 versus 16 percent). Thus, individuals at the left tail of the consumption distribution in the baseline model are both more likely to have stronger net liquid wealth positions (because of their greater patience) and also have a stronger preference for homeownership (because of the smaller size of rental houses they currently reside in). This makes them more likely to become home buyers in the near future, regardless of house-price shocks. When house prices unexpectedly decline this allows some resources, that would have otherwise gone to home purchases, to go to consumption. This effect is accentuated by the fact that borrowing capacity is less restricted in the baseline model compared to the alternative model.

While inequality and poverty are reduced following a decline in house prices, aggregate

[25] In calculating consumption-based poverty in the model we only consider non-housing consumption.

consumption still falls and the decline is larger in the baseline model. The third row of the top panel shows that aggregate consumption falls by 1.8 percent in the baseline model, or by roughly 24 percent more than in the alternative model with no FD. This difference is due to two factors. First, homeownership is higher in the baseline model compared to the alternative model (65 versus 59 percent), making more agents in the former more susceptible to house-price shocks. Second, however, even conditioning on ownership, the MPC out of house-price shocks is larger in the baseline model compared to the alternative model (0.088 versus 0.075).

The final column of this table shows that the previous conclusions are largely unchanged even if we re-estimate the alternative model to match non-FD related moments. The re-estimation process allows the alternative model to better match the homeownership patterns observed across quintiles of FD. As a consequence, the drop in aggregate consumption increases, which attenuates the implied amplification effect of FD. However, the re-estimation process delivers an even more extreme composition of individuals at the left tail of the consumption distribution. As a result, the decrease in consumption inequality in the re-estimated alternative model is smaller, implying an even larger amplification role of FD based on this measure. Overall, though, either alternative model specification implies a similar amplification role of FD, on average. If we take the average (absolute) amplification based on Column (3), we conclude that FD amplifies house-price shocks by 30 percent. The same calculation using Column (4) suggests FD amplifies house-price shocks by 35 percent.

The first and second rows of the bottom panel of Table 8 show that FD also has important consequences for the propagation of labor-income shocks into consumption inequality and poverty. These rows highlight that the baseline model implies a smaller increase in either measure following a labor-income shock when compared to the alternative model with no FD. The numbers in the third column suggest FD reduces the pass-through of labor-income shocks into these measures of consumption by 11 to 30 percent.

The smaller increase in poverty and inequality in the baseline model compared to the alternative model reflects the greater borrowing capacity that individuals at the left tail of the consumption distribution have in this model. For example, individuals below the poverty threshold in the baseline model are, on average, debtors and increase their debt by nearly 14 percent in steady state. In contrast, similar individuals in the alternative model while also debtors, on average, only increase their debt by 1 percent in steady state. If we focus on type-L individuals (to account for compositional differences), we still find large differences in borrowing capacity across models. Type-L individuals below the poverty threshold in the baseline model increase their debt by nearly 12 percent, while their counterparts in the alternative model only increase their debt by 3 percent. Overall, individuals in the baseline model who are in the left tail of the consumption distribution are better able to insulate consumption from temporary labor-income shocks.

The third row of the bottom panel reveals that even though consumption inequality and poverty rise by less in the baseline model compared to the alternative model, aggregate consumption falls a bit more in the former following a labor-income shock. Indeed, aggregate consumption falls by 3.3 percent in the baseline model, nearly 16 percent more compared to the predicted drop in the alternative model. The greater borrowing capacity of individuals in the baseline model means that type-L individuals have a larger consumption share than in the alternative model. And since type-L individuals also have higher MPCs out of labor-income shocks (as shown in Table 7) this means that the response of aggregate consumption in the baseline model will also be larger. Not surprisingly, the overall MPC out of labor-income shocks is indeed about 5 cents larger (30 versus 25 cents) in the baseline versus alternative model.

The last column of the bottom panel again shows that the conclusions pertaining to the amplification or attenuation effect of FD are broadly similar if we instead re-estimate the alternative model. The re-estimation process brings the relationship between consumption and individual ex-ante types in the alternative model more in line with the baseline model. As a consequence, the consumption share of type-L individuals increases and approaches that of the baseline model. This exacerbates the drop in aggregate consumption in the alternative model. As a result, the inferred amplification of labor-income shocks due to FD is reduced somewhat. However, because type-L individuals, and, in particular, those at the left tail of consumption distribution, consume more but still face the same ad hoc borrowing constraint as before, the implied increase in poverty is even larger in the alternative model. This increases the inferred attenuation effect that FD has on this measure. Overall, though, either alternative model specification yields similar implications for the importance of FD, on average. If we take the average absolute values from Column (3), we conclude that FD alters (i.e. either amplifies or attenuates) the response to labor-income shocks by nearly 19 percent. The same calculation using Column (4) suggests FD changes the pass-through of labor-income shocks by 20 percent.

In sum, the results of this section suggest that FD *alters* (i.e. either amplifies or attenuates) the response of various consumption measures by around 25 percent, on average, depending on the type of shock. House-price shocks are affected more (around 30 percent), while labor-income shocks less (around 20 percent). A key question is how much of this conclusion is explained by the correlation of FD with aggregate shocks (that the alternative model lacks) and how much has to do with modeling FD specifically (which the alternative model circumvents with a single ad hoc borrowing constraint).

6.2 How important is the correlation of FD with aggregate shocks?

In the previous section, we demonstrated that including FD alters the responses of aggregate consumption, consumption inequality, and consumption-based poverty to macroeconomic shocks relative to a model with no FD. However, the baseline model imposes the data-consistent correlation structure of aggregate shocks with FD, as in Table 4. In contrast, the alternative model with no FD has, by construction, no relationship between FD and aggregate shocks. In this section, we assess how this correlation structure affects the previous conclusions on the amplification or attenuation effects of FD. Overall, we find that our conclusions are nearly the same even if aggregate shocks are uncorrelated with FD. Thus, our results are mostly driven by differential responses across individuals to the same shock rather than similar people responding to different shocks.

To isolate the importance of the correlation between aggregate shocks and FD we consider a counterfactual baseline economy where shocks are the same across quintiles of FD. Using the numbers from Table 4, we compute the average decline in house prices or earnings. We then hit each of the quintiles with this common shock. As the only difference between this counterfactual economy and our baseline model is the correlation structure of aggregate shocks, differences between these economies help us identify the importance of the correlation structure of aggregate shocks in driving our main conclusions.

Table 9 summarizes the results of this exercise and suggests that our main conclusions are largely independent of the correlation structure of aggregate shocks with FD. The first column of this table presents the percent difference in the corresponding consumption measure between the baseline model and the counterfactual model with uncorrelated shocks, as a percent of the baseline model's prediction (i.e. as a fraction of Column (1) from Table 8). Since Column (1) is presented as a fraction of the baseline model's prediction it can be interpreted as the contribution of the correlation of aggregate shocks with FD to the amplification/attenuation effects obtained in the previous section. Consistently, Columns (2) and (3) present the remaining percentage of the amplification/attenuation effect that is due to FD alone. Column (2) presents the remainder using the amplification/attenuation effect implied from Column (3) of Table 8 (i.e. when the alternative model with no FD is *not* re-estimated), while Column (3) presents the remainder using the amplification/attenuation effect implied from Column (4) (i.e. when the alternative model with no FD *is* re-estimated).

Table 9: The importance of the correlation of FD with aggregate shocks by shock and consumption measure

	Correlation of FD w/shocks (1)	FD alone (2)	FD alone (re-estimated) (3)
House-price shocks			
Consumption p90/p10	5.09	14.16	45.22
Consumption-based poverty	-10.96	58.04	55.52
Aggregate consumption	1.89	22.60	8.45
Average absolute value		31.60	36.40
Labor-income shocks			
Consumption p90/p10	9.10	-19.99	-14.44
Consumption-based poverty	6.52	-36.68	-52.12
Aggregate consumption	1.51	14.34	6.75
Average absolute value		23.67	24.44

Notes: Column (1) presents the percent difference in the corresponding measure between the baseline model with correlated (Column 1 from Table 8) and uncorrelated shocks, as a fraction of the former. Column (2) presents the percent difference in the corresponding measure between between the baseline model with uncorrelated shocks and the alternative model with no FD (Column 2 from Table 8), as a fraction of the former. Column (3) presents the percent difference in the corresponding measure between the baseline model with uncorrelated shocks and the re-estimated alternative model with no FD, as a fraction of the former. The average absolute value represents the average of the absolute values of either Column (2) or (3).

Focusing on the top panel of Table 9 reveals that the correlation of aggregate shocks with FD plays a fairly minor role in FD altering the pass-through of house-price shocks into the various measures of consumption. An easy way to see this is by looking at the average absolute value in Column (2) and comparing it to its counterpart in Column (3) of Table 8. The value of 31.6 percent in this table suggests that all of the amplification effect (e.g. 30.28 percent) of house-price shocks measured in Table 8 is accounted for by FD alone. Column (3) of this table shows that this conclusion is unchanged when the alternative model (which we use as the yardstick to measure the overall amplification effect of FD) is re-estimated. Here again, the average absolute value of the amplification effect is fully accounted for by FD alone.

Turning to the bottom panel of Table 9 suggests the correlation of aggregate shocks plays a slightly larger, though still modest role, on average, in FD altering the pass-through of labor-earnings shocks into the various measures of consumption. Looking at the average absolute values in Columns (3) and (4) shows that removing the correlation of aggregate shocks with FD increases the average amplification effect of FD for labor-income shocks. However, the additional amplification is still modest. Taking Column (3) as an example, removing the correlation of aggregate labor-income shocks with FD increases the average absolute amplification by 4.7 percent.

The only case where the correlation of FD with aggregate shocks matters significantly is when looking at the response of consumption inequality to labor-income shocks. Excluding the correlation of FD with aggregate labor-income shocks nearly doubles the attenuation effect of FD from -11 percent (Column (3) in Table 8) to -20 percent (Column (2) of Table 9). This observation is easily explained by the quintile provenance of individuals at the left and right tails of the consumption distribution in the steady state of the baseline model. The median individual in the left tail of the consumption distribution is from Q4 of the FD distribution (i.e. faces larger shocks than average). Thus, removing the correlation of aggregate shocks with FD reduces the size of the shock faced by many of those at the bottom of the consumption distribution. Similarly, the median individual at the top of the consumption distribution is from Q2 (i.e. faces smaller shocks than average). Thus, removing the correlation of aggregate shocks with FD increases the size of the shock faced by many of those at the top of the consumption distribution. These two forces tend to reduce the increase in consumption inequality in the baseline model with FD, which to begin with is already smaller than in the alternative model with no FD. Overall, removing the correlation of shocks with FD increases the attenuation effect that FD has on the pass-through of labor-income shocks into consumption inequality.

6.3 FD and the housing market

Having established how FD alters the transmission of shocks into consumption, we now explore how allowing for FD affects allocations in the housing market. To do so, we compare a few key statistics from the steady states of the baseline model with FD and the alternative model with no FD and only a borrowing constraint. Here we only consider the alternative model which isn't re-estimated as we want to keep preferences constant across models and only change the debt market arrangement.

As can be seen from Table 10, FD has consequences for the housing market that are distinct from a model with only an ad hoc borrowing constraint. While the top row of this table shows some differences in homeownership across the two models, the more noteworthy divergences appear in the subsequent rows. The second row shows that housing leverage is roughly twice as high (regardless of quintile) in the baseline model with FD as it is in the alternative model with an ad-hoc borrowing constraint for unsecured credit. The third and fourth rows show that the fraction of those having housing debt and the ratio of housing debt to income are higher in the model with FD than in the alternative model. The last row of this table demonstrates that mortgage default rates are substantially greater in the model with FD as a result of the increased indebtedness. Except for Q1 (least distressed), all quintiles show mortgage default rates that are about twice as high when FD is permitted as when it is not.

Table 10: The housing market implications of FD

	Q1		Q2		Q3		Q4		Q5	
	Baseline model	Alt. model	Baseline model	Alt model	Baseline model	Alt model	Baseline model	Alt model	Baseline model	Alt. model
Home ownership*	76.63	70.54	71.04	64.56	62.34	56.43	61.63	55.08	52.58	46.10
Housing leverage	30.06	15.61	36.56	19.00	40.89	18.76	44.34	22.70	44.74	22.21
Housing debt> 0*	28.75	21.82	28.72	21.46	25.04	17.01	27.43	19.72	24.2	15.96
Housing debt/Inc	0.86	0.54	1.03	0.64	1.17	0.64	1.31	0.80	1.52	0.91
Mortg def rate*	1.3	1.08	1.69	0.66	2.2	1.13	2.18	0.76	2.49	0.86

Notes: $*$ in percent.

Echoing the previous sections, these findings are largely explained by the differing composition across models of who owns homes. In the baseline model, roughly 15 percent of homeowners are type-L individuals. In contrast, in the alternative model, only

7 percent of homeowners are type-L individuals. As type-L homeowners have worse financial asset positions and larger mortgages compared to their type-H counterparts, the pool of homeowners in the baseline model is naturally more likely to default on their mortgages. Indeed, regardless of the model, the mortgage default rate among type-L homeowners is around 4 percent, whereas the default rate among type-H owners is under 1 percent. It is important to recall that preferences across both models are the same. Thus, these results largely reflect how the composition of homeownership changes when the unsecured debt market structure is altered.

Why are type-L individuals more likely to own homes in the baseline model with FD? This is largely are reflection of the interaction between FD in the unsecured debt market and homeownership. When borrowing rates are endogenous (as in the baseline model with FD), creditors account for the fact that homeowners can extract equity (via refinancing), sell their homes, or default on mortgage obligations instead of defaulting on unsecured credit. In contrast, renters have none of these options. As a result, homeowners (and homebuyers) are offered better terms of credit relative to renters. This additional benefit of homeownership does not exist in the alternative model without FD. Without FD borrowing rates are fixed and do not depend on homeownership status.

7 Concluding remarks

Our paper aims to understand how household financial distress (FD) shapes individual and aggregate consumption dynamics after aggregate shocks. FD, defined as whether households are over 30 days delinquent on paying back unsecured debt, can affect consumption because its presence reflects weak balance-sheet health, and thus differential access to credit, which is critical when buffering consumption from shocks. However, FD may also affect consumption because, as we documented, exposure to aggregate shocks was correlated with prior FD over the past two recessions.

We find that FD *alters* (i.e. either amplifies or attenuates) the response of the aggregate consumption, consumption inequality, and consumption-based poverty by nearly 25 percent, on average, relative to models that omit FD and only allow borrowing up to a fixed limit. House price shocks are altered more (around 30 percent), while labor-income shocks are less (around 20 percent). That prior FD was correlated with aggregate shock exposure matters less for driving these results. Rather, FD matters both because it allows for an additional margin of adjustment when shocks arrive (debt repudiation), but also because matching it implies significant ex-ante heterogeneity in the population, which translates into wide dispersion in consumption responses (e.g. MPCs).

Beyond consumption, we also find that modeling FD has important consequences for homeownership. Unlike a model without FD, in a model with FD, unsecured debt prices are endogenous and adjust with ownership status. In particular, debt prices account for the fact that homeowners can extract equity or default on mortgage payments to pay for debt. As a result, borrowing terms are better for homeowners, and this provides another benefit of homeownership relative to renting. This effect encourages some additional ownership among more financially fragile individuals that otherwise remain as renters in a model without FD. Because these individuals indeed are more likely to default on mortgage payments, overall mortgage default rates are higher in a model with FD compared to one without it.

In terms of future research, our model is very well suited for assessing the impacts of emergency policies enacted during the COVID-19 pandemic, like mortgage and debt forbearance. Additionally, our model can be used to understand survey results in Coibion, Gorodnichenko, and Weber (2020) showing that most households used stimulus checks to pay down debt and improve their financial positions. We intend to address these exciting questions in future work.

References

Aguiar, M., Bils, M., and Boar, C. Who are the hand-to-mouth?, October 2020.

Aiyagari, R. Uninsured idiosyncratic risk and aggregate saving. *The Quarterly Journal of Economics*, 109(4):659–684, 1994.

Aladangady, A. Housing wealth and consumption: Evidence from geographically-linked microdata. *American Economic Review*, 107(11):3415–46, 2017.

Armstrong, G., Cho, C., Garner, T. I., Matsumoto, B., Munoz, J., and Schild, J. Building a consumption poverty measure: Initial results following recommendations of a federal interagency working group. *AEA Papers and Proceedings*, 112:335–39, May 2022.

Aruoba, B., Elul, R., and Kalemli-Ozcan, S. Housing wealth and consumption: The role of heterogeneous credit constraints. 2018.

Athreya, K. Welfare implications of the bankruptcy reform act of 1999. *Journal of Monetary Economics*, 49:1567–1595, 2002.

Athreya, K., Sánchez, J. M., Tam, X. S., and Young, E. R. Labor market upheaval, default regulations, and consumer debt. *Review of Economic Dynamics*, 18(1):32–52, January 2015.

Athreya, K., Sánchez, J. M., Tam, X. S., and Young, E. R. Bankruptcy and delinquency in a model of unsecured debt. *International Economic Review*, 59(2):593–623, 2017.

Athreya, K., Mustre-del Río, J., and Sánchez, J. M. The persistence of financial distress. *The Review of Financial Studies*, 32(10):3851–3883, 2019.

Auclert, A. and Mitman, K. Consumer bankruptcy as aggregate demand management. 2019.

Becker, S. G. and Mulligan, C. B. The endogenous determination of time preference. *The Quarterly Journal of Economics*, 112(3):729–758, August 1997.

Berger, D. and Vavra, J. Shocks versus Responsiveness: What Drives Time-Varying Dispersion? *Journal of Political Economy*, 127(5):21042042, 2019.

Berger, D., Guerrieri, V., Lorenzoni, G., and Vavra, J. House Prices and Consumer Spending. *Review of Economic Studies*, 85(3):1502–1542, 2018.

Bick, A. and Blandin, A. Real-time labor market estimates during the 2020 coronavirus outbreak. February 2021.

Calvet, L. E., Campbell, J. Y., Gomes, F. J., and Sodini, P. The cross-section of household preferences, 2019.

Chatterjee, S., Corbae, D., Nakajima, M., and Ríos-Rull, J.-V. A quantitative theory of unsecured consumer credit with risk of default. *Econometrica*, 75(6):1525–1591, June 2007.

Coibion, O., Gorodnichenko, Y., and Weber, M. How Did U.S. Consumers Use Their Stimulus Payments? NBER Working Papers 27384, National Bureau of Economic Research, Inc, 2020.

Corbae, D. and Quintin, E. Leverage and the Foreclosure Crisis. *Journal of Political Economy*, 123(1):1–65, 2015.

Coronado, J. L., Lupton, J. P., and Sheiner, L. The household spending response to the 2003 tax cut: evidence from survey data. Technical report, 2005.

Dupor, B., Mehkari, M. S., Li, R., and Tsai, Y.-C. The 2008 US Auto Market Collapse. Technical report, 2019.

Garriga, C. and Hedlund, A. Mortgage Debt, Consumption, and Illiquid Housing Markets in the Great Recession. October 2017.

Gregory, V., Menzio, G., and Wiczer, D. The alpha beta gamma of the labor market. NBER Working Paper No. 28663, April 2021.

Gruber, J. and Martin, R. Precautionary savings and the wealth distribution with illiquid durables. 2003.

Hanushek, E. A. and Quigley, J. M. What is the Price Elasticity of Housing Demand? *The Review of Economics and Statistics,*, 62(3):449–454, August 1980.

Hatchondo, J. C., Martinez, L., and Sánchez, J. M. Mortgage defaults. *Journal of Monetary Economics*, 76:173–190, November 2015.

Herkenhoff, K. F. The role of default (not bankruptcy) as unemployment insurance: New facts and theory. Working Paper, January 2013.

Herkenhoff, K. F. and Ohanian, L. E. Foreclosure delay and u.s. unemployment. Working Paper, March 2012.

Huggett, M. The risk-free rate in heterogeneous-agent incomplete-insurance economies. *Journal of Economic Dynamics and Control*, 17:953–969, 1993.

Jappelli, T. and Pistaferri, L. Intertemporal Choice and Consumption Mobility. *Journal of the European Economic Association*, 4(1):75–115, March 2006.

Jeske, K., Krueger, D., and Mitman, K. Housing, mortgage bailout guarantees and the macro economy. *Journal of Monetary Economics*, 60(8):917–935, November 2013.

Kaplan, G. and Violante, G. L. How Much Consumption Insurance beyond Self-Insurance? *American Economic Journal: Macroeconomics*, 2(4):53–87, October 2010.

Kaplan, G. and Violante, G. L. A model of the consumption response to fiscal stimulus payments. *Econometrica*, 82(4):1199–1239, July 2014.

Kaplan, G., Mitman, K., and Violante, G. L. Non-durable Consumption and Housing Net Worth in the Great Recession: Evidence from Easily Accessible Data. NBER Working Papers 22232, National Bureau of Economic Research, Inc, May 2016.

Kaplan, G., Moll, B., and Violante, G. L. Monetary Policy According to HANK. *American Economic Review*, 108(3):697–743, March 2018.

Kaplan, G., Mitman, K., and Violante, G. L. The housing boom and bust: Model meets evidence. *Journal of Political Economy*, 128(9):3286–3345, 2020.

Krusell, P. and Smith., A. A. Income and Wealth Heterogeneity in the Macroeconomy. *Journal of Political Economy*, 106(5):867–896, October 1998.

Kuhn, M. and Ríos-Rull, J.-V. 2013 update on the u.s. earnings, income, and wealth distributional facts: A view from macroeconomics. *Federal Reserve Bank of Minneapolis Quarterly Review*, 37(1):1–73, April 2016.

Li, W., Liu, H., Yang, F., and Yao, R. Housing over Time and over the Life Cycle: A Structural Estimation. Working paper no. 15-04, Federal Reserve Bank of Philadelphia, 2015.

Livshits, I., MacGee, J., and Tertilt, M. Consumer bankruptcy: A fresh start. *American Economic Review*, 97(1):402–18, March 2007.

Meyer, B. D. and Sullivan, J. X. Annual report on u.s. consumption poverty: 2018. Manuscript, October 2019.

Mian, A., Rao, K., and Sufi, A. Household Balance Sheets, Consumption, and the Economic Slump. *The Quarterly Journal of Economics*, 128(4):1687–1726, 2013.

Mustre-del Río, J. Wealth and labor supply heterogeneity. *Review of Economic Dynamics*, 18:619–634, 2015.

Parker, J. A. Why don't households smooth consumption? evidence from a $25 million experiment. *American Economic Journal: Macroeconomics*, 9(4):153–183, 2017.

Pennington-Cross, A. The duration of foreclosures in the subprime mortgage market: A competing risks model with mixing. Working Paper 027A, Federal Reserve Bank of St. Louis, 2006.

Sahm, C. R., Shapiro, M. D., and Slemrod, J. Household Response to the 2008 Tax Rebate: Survey Evidence and Aggregate Implications. In *Tax Policy and the Economy, Volume 24*, NBER Chapters, pages 69–110. National Bureau of Economic Research, Inc, 2010.

Saiz, A. The Determinants of Housing Supply. *The Quarterly Journal of Economics*, 125 (3):1253–1296, 2010.

Siegel, S. Consumption-Based Asset Pricing: Durable Goods, Adjustment Costs, and Aggregation. Working paper, University of Washington, 2008.

Online Appendix

A Empirical analysis

In the following subsections, we present detailed information about each variable and how it was constructed, as well as various empirical results to supplement those shown in the paper. Table A1 shows some initial summary statistics for the entire data set. The next subsections explain how the data set was constructed.

Table A1: Descriptive Statistics Across Zip codes

	Count	Mean	S.D.	p25	p50	p75
Housing Net Worth Shock, 2006-9	14230	-0.098	1.035	-0.109	-0.030	0.005
Change in home value \$000, 2006-9	14230	-38.905	64.130	-62.833	-13.200	2.300
Net Worth per Household \$000, 2006	14230	487.854	934.963	159.956	269.338	496.700
Income Per Households, \$000, 2006	14230	72.861	53.508	45.125	58.838	82.823
No. Hou. per zip code (ths), 2006	14230	11.390	6.399	6.703	10.968	15.305
Housing Leverage Ratio, 2006	14230	0.453	0.173	0.347	0.433	0.536
Δ_{06-09} auto spending per hou. \$000	14230	-2.108	6.447	-2.525	-1.517	-0.835
Fraction in DQ30, 2006	14230	0.142	0.048	0.108	0.138	0.172
Fraction in CL80, 2006	14230	0.228	0.054	0.192	0.228	0.264

Notes: All statistics are weighted by the number of households in the first quarter of 2006 for each zip code. p25, p50, and p75 respectively give the 25th, 50th, and 75th percentiles.
Sources: IRS SOI, FRBNY Consumer Credit Panel/Equifax, Census Bureau, Zillow, SCF.

A.1 A geographically representative sample

Building a geographically representative sample from the FRBNY CCP/Equifax dataset over all the years considered in this study presents a slight challenge: small random samples will give good estimates at the national level, and even for the largest zip codes, but poor estimates for the smallest zip codes. Using much larger random samples over the full country could fix this issue, but the resulting datasets become difficult to process. Instead, then, we divide the zip codes for which we have IRS Summary of Income (SOI) data into 10 groups by population size[26] and oversample areas with lower population.

Specifically, we pull a 100 percent sample of individual Equifax records from the smallest zip codes by population and decrease that percentage linearly until we pull a 50 percent sample of Equifax records for the largest zip codes.[27] To remain in our sample for a given quarter, individuals must be between 25 and 65 years old, inclusive.[28] Then, we

[26] Specifically by using the “number of returns” field provided by the IRS SOI.

[27] Zip-code level data on CL80 and DQ30 are available at this link for the years 2006 and 2018.

[28] Age is calculated using an individual’s recorded birth year, and so any records not including a birth year are also excluded.

correct for oversampling by reweighting using population data from the 2000 and 2010 Census.

A.2 Constructing measures of wealth and consumption

The household wealth portion of our dataset was constructed at the zip code and county levels using a method almost identical to that of Mian, Rao, and Sufi (2013). Net wealth is defined as the sum of housing wealth H and financial wealth FW less debt D. H is calculated separately for zip codes and counties as the median home value multiplied by the number of owner-occupied housing units in each geography. We use Zillow data for home values and Census data on owner-occupied housing units.[29] The housing leverage ratio is then defined as the total housing debt in a geography divided by H. Total housing debt is the mean housing debt[30] recorded in Equifax for each geography multiplied by the number of households in that geography, taken from the Census.

To construct FW, we began by using IRS SOI data to calculate the fraction of national interest and dividends held by a given zip code. Then, each zip code was apportioned a share of the national financial wealth recorded in the Survey of Consumer Finances (SCF) corresponding to that fraction.[31] FW at the county level is simply calculated as the sum of FW in its component zip codes.[32] D is calculated similarly to FW. First, we calculate the fraction of the total debt balance in our sample of the Equifax dataset accounted for by a given zip code or county.[33] Next, we assign each geography a portion of the total debt from the SCF equal to that fraction.

[29]To fill in the missing years in Census data, we interpolate owner-occupied housing units linearly for each zip code and county from 2000 to 2010. Mian, Rao, and Sufi (2013) did not use Zillow data for home values and instead relied entirely on home price information from the 2000 Census tracked upward through time by the Core Logic price index. Using Zillow data affords us the advantage of much wider data coverage.

[30]Here we include mortgages, the home equity installment balance, and the home equity revolving balance.

[31]Mian, Rao, and Sufi (2013) used the Federal Flow of Funds for this purpose, but we use the Survey of Consumer Finances because it allows us to limit our financial wealth totals to those of a certain age range. Specifically, our model is calibrated to match dynamics among people who are 25 to 55 years old, and so we likewise restrict the data to that age range when setting calibration targets. As shown in Kuhn and Ríos-Rull (2016), the SCF and Federal Flow of Funds match up quite nicely in terms of aggregates. The SCF is not available in every year, and so wherever necessary we interpolate linearly between available years.

[32]To avoid double counting FW, this requires that something be done about zip codes that span multiple counties. We elected to assign all of a zip code's FW into the county that most people in that zip code inhabit.

[33]Because our method of pulling Equifax data intentionally oversampled geographic areas with lower populations, we weight each geography's debt by the number of households it encompasses in the Census.

In the regression analysis, we follow Mian, Rao, and Sufi (2013) in using new car registrations as a measurement of consumption. Specifically, we use data from R.L. Polk by IHS Markit to find the number of new automobiles registered each year by residents of each zip code and county. As noted by Mian, Rao, and Sufi (2013), these data are advantageous relative to other sources of consumption data because they record where the car buyer lives rather than the point of sale, but disadvantageous in that they do not include the price of each vehicle purchased. To resolve this issue, we use the method Mian, Rao, and Sufi (2013) in allocating an annual share of the national Census Retail Trade amounts for "Auto, Other Motor Vehicle" to each zip code and county equal to the share of new autos that residents of each geography purchased in the Polk data. Recall that the main interest of our regression analysis is to evaluate whether zip code-level MPCs vary with the level of FD within a zip code, which motivates our link between FD and vulnerability.

A.3 Financial distress and the house-price Shock

As defined in Section 2, DQ30 gives the percentage of primary borrowers in the Equifax dataset who are at least 30 days delinquent on a credit card payment during some quarter of the year. CL80 was similarly defined for primary borrowers as the percentage of people who have reached at least 80 percent of their credit limit during some quarter of the year.

First, we show that the correlation we established in Figure 3 in the main text holds if we replace home values with housing wealth shocks as in Mian, Rao, and Sufi (2013). Figure A1 documents the main result: The incidence of the housing wealth shock upon zip codes was highly positively correlated with household FD so that zip codes with higher FD in 2002 tended to experience heavier losses in the Great Recession.

Then, we show this correlation is robust to alternative definitions of FD, as can be seen in Figure A2. The levels of FD change depending on the definition, but the corresponding pattern in the housing net worth shock is immediately apparent in every case.

As would be expected from the persistence of FD, these results are also not dependent upon measuring FD in a particular year. Figure A3 shows that the same relationship holds when measuring FD just before the recession started in 2006.

Figure A1: Housing Wealth Shocks (2006-09) and FD (DQ30) in 2002

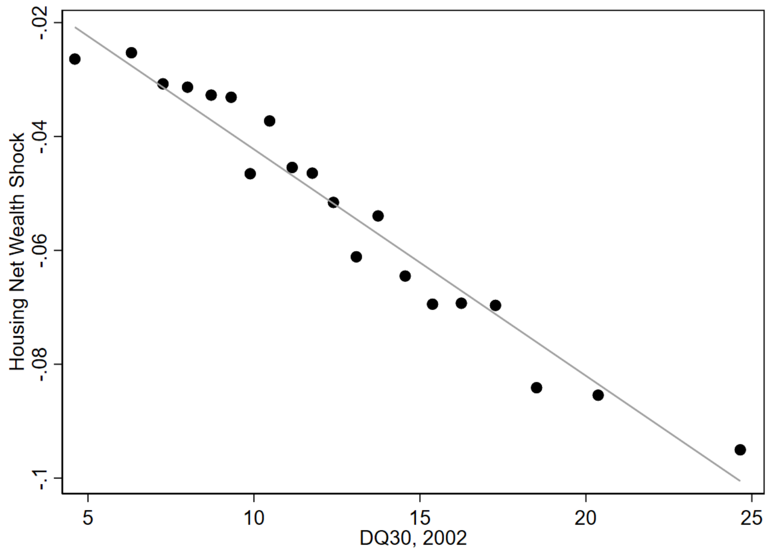

Sources: IRS SOI, Zillow, FRBNY Consumer Credit Panel/Equifax, Census Bureau, SCF. Each dot represents the mean of that bin weighted by the 2006 net wealth of bins with respect to DQ30.

Figure A2: Robustness to the Definition of FD

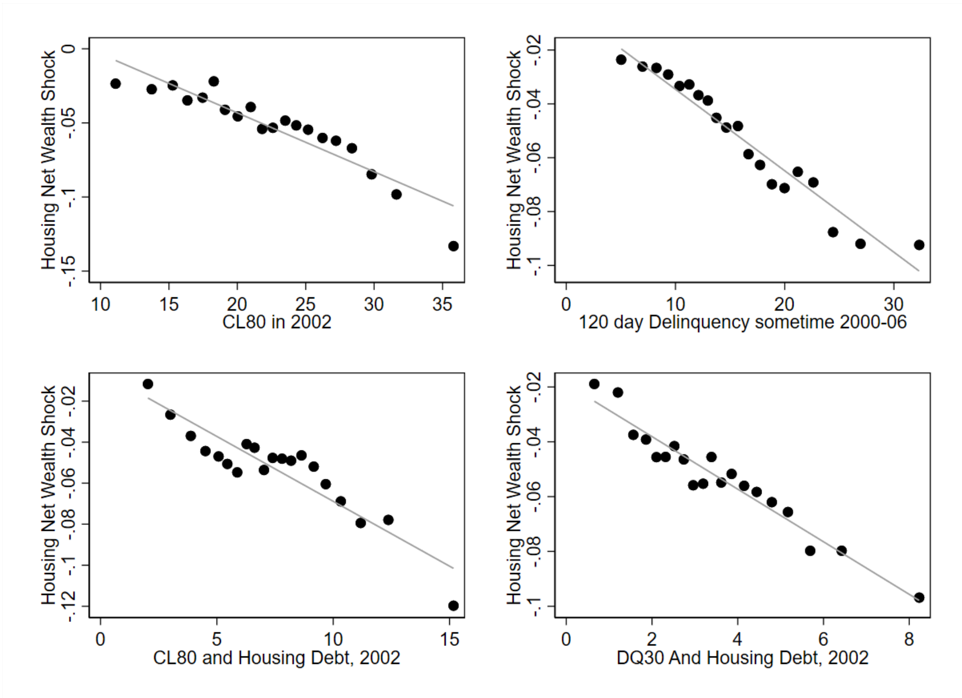

Notes: "120-day Delinquency sometime 2000-06" gives the percent of people in a zip code who were 120 days or more delinquent on credit card payments at least once between 2000 and 2006. "CL80 and Housing Debt, 2002" gives the percentage of people in a zip code both in FD under the CL80 definition and having debt indicative of owning a house (i.e., a mortgage or home equity line of credit). "DQ30 and Housing Debt, 2002" is similar.
Sources: IRS SOI, Zillow, FRBNY Consumer Credit Panel/Equifax, Census Bureau, SCF. Each dot represents the mean of that bin of FD weighted by 2006 net wealth.

Figure A3: Robustness to Measuring FD in 2006

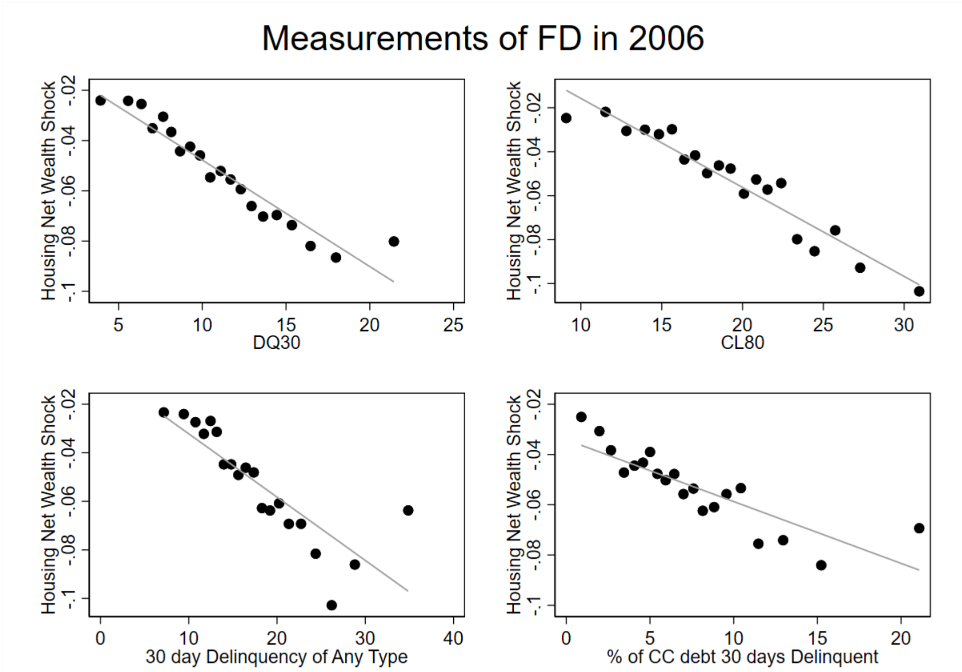

Notes: "30-day Delinquency of Any Type" gives the percentage of people in a zip code that are 30 or more days delinquent on any type of debt as recorded by the New York Federal Reserve Bank/Equifax CCP. "% of CC debt 30 days Delinquent" gives the percentage of all credit card debt in a zip code that is at least 30 days delinquent.
Sources: IRS SOI, Zillow, FRBNY Consumer Credit Panel/Equifax, Census Bureau, SCF. Each dot represents the mean of that bin of FD weighted by 2006 net wealth.

A.4 Regressions

There is an increasing relationship between FD and a zip code's marginal propensity to consume, as illustrated in Figure 2: The more prevalent FD within a zip code, the more its residents tended to cut consumption of autos in response to a dollar decline in their housing wealth during the Great Recession. This section presents the regression results used to construct that figure, further motivating the importance our model ascribes to FD in shaping consumption patterns both for individual regions and the aggregate economy.

Table A2 reports the baseline results. In addition to the usual measurements of FD, we include two additional metrics for robustness: "DQ30 and CL80" calculates for each individual the portion of quarters in a year that they spent with either a credit card payment 30 days delinquent or having reached 80 percent of their credit limit[34] and then

[34]To give a clarifying example, say that there was an individual who in quarter 1 of 2002 was both at least 30 days delinquent on a credit card payment and had used over 80 percent of their available credit card limit. Then, in quarter 2, they remained over 80 percent of their credit card limit but did not have any credit card payments over 30 days delinquent. The rest of the year occurred without any credit

averages that percentage across the geography. "ADQ30" is defined much like DQ30, but gives the percentage of people in a zip code who are at least 30 days delinquent on any kind of debt recorded by the FRBNY/Equifax CCP. All columns reveal statistically significant coefficients at the 0.001 level for house price shocks (i.e., the change in home value between 2006 and 2009) and the interaction of these shocks with FD. Comparing across columns suggests that our estimated coefficients are robust to the definition of FD we use. Importantly, the interaction term is positive: higher FD in 2002 is associated with larger consumption drops between 2006 and 2009.

Table A2: Auto Spending at the zip-code level

	Δ_{06-09} Auto Spending			
FD Measure in 2002:	(DQ30)	(CL80)	(CL80 and DQ30)	(ADQ30)
Δ_{06-09} Home Value	-0.005	-0.008	-0.009*	-0.006
	(0.00)	(0.00)	(0.00)	(0.00)
FD	-5.635***	-5.479***	-5.893***	-3.893***
	(1.16)	(1.03)	(1.20)	(0.75)
Δ_{06-09} Home Value $\times$ FD	0.101***	0.071***	0.100***	0.072***
	(0.02)	(0.02)	(0.02)	(0.01)
Observations	14038	14038	14038	14038

Notes: Controls include change in income and change in financial wealth and the interaction of these variables with the alternative variables of FD. We additionally control for the interactions between changes in home values, financial wealth, and income and the 2006 levels of income and financial wealth, as well as the percent of households that owned homes in 2006 and a constant. All regressions are weighted by the number of owner-occupied housing units in the zip code as of 2006. Standard errors appear in parentheses.

Given the results of Mian, Rao, and Sufi (2013), it may be worried that FD in these regressions is merely capturing variation in housing leverage. Figure 2 directly compares our baseline to the results controlling for the housing leverage ratio and Table A3 shows the corresponding regression output. The results for the interaction term of interest remain near unchanged for every measure of FD included, removing these concerns. Indeed, as shown in Figure A4, there is if anything a negative relationship between FD measured in 2002 and housing leverage in 2006; i.e., regions with more financial distress in 2002 tend to have lower leverage in 2006. The other panel shows that there does not appear to be as clear of a contemporaneous relationship between FD and housing leverage in 2002.

There may also be a concern that our measurement of FD is really capturing dif-

incident. On our metric, this individual would have spent 50 percent of the year in financial distress. Similar calculations would be made for all other individuals in our sample from their geography, and those numbers would be averaged to reach the final result.

Table A3: Auto Spending at the Zip-code Level Controlling for Leverage

	Δ_{06-09} Auto Spending			
FD Measurement taken in 2002:	(DQ30)	(CL80)	(DQ30 and CL80)	(ADQ30)
Δ_{06-09} Home Value	-0.012*	-0.013*	-0.015*	-0.013*
	(0.01)	(0.01)	(0.01)	(0.01)
FD	-5.458	-7.239*	-7.495*	-4.548*
	(3.13)	(2.89)	(3.32)	(2.02)
Δ_{06-09} Home Value × FD	0.104***	0.068***	0.097***	0.073***
	(0.02)	(0.02)	(0.02)	(0.01)
Housing Leverage Ratio$_{06}$	-0.228	-1.216	-0.953	-0.677
	(1.15)	(1.69)	(1.48)	(1.18)
Δ_{06-09} Home Value × Housing Leverage Ratio$_{06}$	0.018*	0.014	0.016*	0.019*
	(0.01)	(0.01)	(0.01)	(0.01)
Housing Leverage Ratio, 2006 × FD	-0.320	4.519	4.164	1.637
	(6.69)	(6.16)	(7.11)	(4.37)
Observations	14136	14136	14136	14136

Notes: Regressions are weighted by the number of owner-occupied housing units in each county in 2006. Additional controls not shown here include the change in income; the change in financial wealth; and interactions between changes and levels for income, financial wealth, and housing wealth. The changes in income and financial wealth are also interacted with leverage.

Figure A4: Correlation of Housing Leverage (2002 and 2006) with FD (DQ30) in 2002

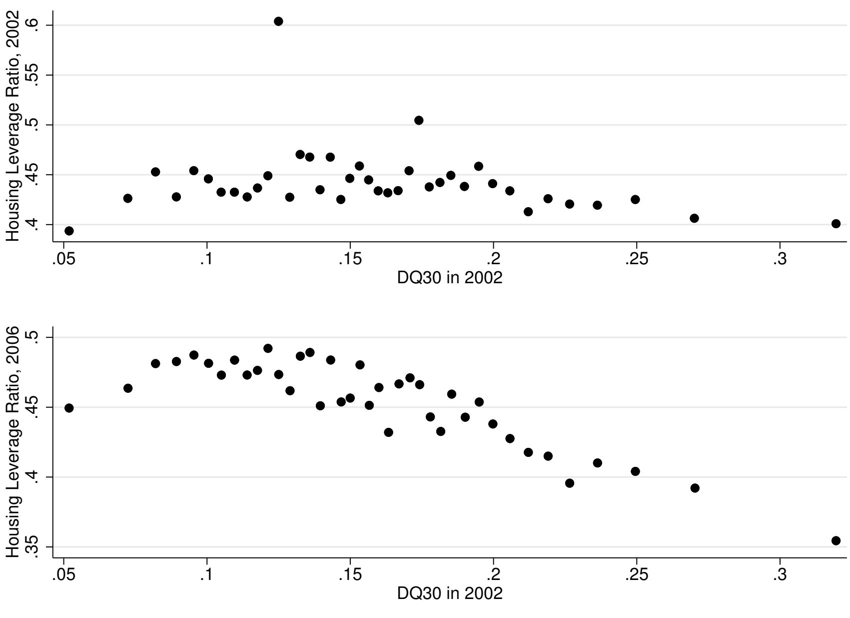

Notes: Housing leverage is here measured as housing debt (including mortgages and home equity lines of credit) divided by the total housing wealth in each geography. For ease of viewing, the data have been divided into 40 bins with respect to CL80, and each dot represents the mean of that bin weighted by the number of households in each zip code as of 2006.

ferences across zip codes in terms of income volatility, where areas with more FD have greater volatility. However, table A4 shows that the opposite is actually true, where quintiles with higher FD actually tend to have less income volatility. Here, we measure income volatility as the standard deviation of the year-over-year percent change in zip code adjusted gross income per household for the years 1999-2005. Table A5 shows that including this additional variable does not remove the effect of the interaction between changing home values and FD shown in table A2. In fact, these effects become slightly larger when the new controls are added. Tables A4 and A5 also include information on the share of zip code employment in manufacturing, which is an industry with output that can be traded across regions. This share is non-monotonic across quintiles of FD and does not remove the effect of the interaction between changing home values and FD.

	FD 2002 Quintile				
Measure	Q1	Q2	Q3	Q4	Q5
DQ30 Income volatility (adj)	1.06	0.92	0.83	0.84	0.79
DQ30 Manufacturing %	11.58	12.36	12.35	12.06	10.72
DQ30 Service %	19.74	21.06	21.62	22.20	22.31
ADQ30 Income volatility (adj)	1.05	0.95	0.83	0.86	0.75
ADQ30 Manufacturing %	10.91	11.60	12.38	12.61	11.59
ADQ30 Service %	19.68	21.15	21.71	22.25	22.13
CL80 Income volatility (adj)	1.04	0.88	0.88	0.82	0.82
CL80 Manufacturing %	11.50	12.62	12.56	12.00	10.41
CL80 Service %	19.64	21.02	21.70	22.11	22.43

Table A4: Zip code variables - Quintile Averages (2002 FD)

Table A5: Auto spending at the zip-code level with controls for Income Volatility and some Industry Shares

	Δ_{06-09} Auto Spending			
FD Measure in 2002:	(DQ30)	(CL80)	(CL80 and DQ30)	(ADQ30)
Δ_{06-09} Home Value	-0.009	-0.011	-0.013	-0.009
	(0.01)	(0.01)	(0.01)	(0.01)
FD	-4.987***	-4.822***	-5.127***	-3.564***
	(1.19)	(1.06)	(1.24)	(0.78)
Δ_{06-09} Home Value × FD	0.098***	0.070***	0.099***	0.070***
	(0.02)	(0.02)	(0.02)	(0.01)
Income Volatility	-0.025*	-0.025*	-0.026*	-0.025*
	(0.01)	(0.01)	(0.01)	(0.01)
Percent employed in manufacturing industry 2005	0.015	0.015	0.016	0.019*
	(0.01)	(0.01)	(0.01)	(0.01)
Percent employed in service industry 2005	-0.039*	-0.035*	-0.036*	-0.035*
	(0.02)	(0.02)	(0.02)	(0.02)
Δ_06-09 Home Value × Manufacturing Share	0.000	0.000	0.000	0.000
	(0.00)	(0.00)	(0.00)	(0.00)
Δ_{06_09} Home Value × Service Share	0.000	0.000	0.000	0.000
	(0.00)	(0.00)	(0.00)	(0.00)
Observations	14038	14038	14038	14038

Notes: All of the controls from table A2 are included here as well. All regressions are weighted by the number of owner-occupied housing units in the zip code as of 2006. Standard errors appear in parentheses.

If we aggregate to the county level, we can also control for the unemployment rate in 2009 and the percentage of employment in tradable industries. Table A6 shows that our main results are robust to these controls as well.

Table A6: Auto Spending at the County Level

	Δ_{06-09} Auto Spending			
FD Measure in 2002	(DQ30)	(CL80)	(DQ30 and CL80)	(ADQ30)
Δ_{06-09} Home Value	-0.090	-0.144	-0.111	-0.088
	(0.07)	(0.09)	(0.09)	(0.07)
FD	-21.165	-10.369	-17.695	-10.406
	(14.15)	(14.35)	(16.13)	(8.61)
Δ_{06-09} Home Value × FD	0.861**	0.679**	0.774**	0.563**
	(0.28)	(0.26)	(0.30)	(0.19)
Income Volatility	-0.319**	-0.292*	-0.293*	-0.299*
	(0.12)	(0.12)	(0.12)	(0.12)
Unemployment Rate, 2009	-0.086	-0.150	-0.108	-0.063
	(0.23)	(0.23)	(0.23)	(0.23)
% of Employment in Tradable Industries	0.092	0.134	0.124	0.113
	(0.08)	(0.08)	(0.08)	(0.08)
Observations	659	659	659	659

Notes: All of the controls from table A2 are included here as well, but aggregated to the county level. All regressions are weighted by the number of owner-occupied housing units in the county as of 2006. Standard errors appear in parentheses.

To mitigate the risk that their results stem from an omitted variable correlated with the decline in home prices, Mian, Rao, and Sufi (2013) instrument for changes in home value using housing supply elasticities from Saiz (2010). Our results are robust to these considerations as well, as shown in table A7, where we present the second stage of a regression that instruments for the change in housing prices. Since we are allowing here the possibility that omitted variable bias is affecting results on the change in home prices, we also need to instrument for the interactions between this change in home prices and other variables. This requires a separate first-stage calculation for each of those variables, and so to simplify things we remove the interactions between the change in housing wealth and the 2006 levels of income and financial wealth. We instrument for the interaction between the change in home prices and FD using the interaction between FD and the housing supply elasticity. While there is not sufficient power for the coefficients on this term to be statistically significant, the estimated values are closely in line with the estimates from previous tables.

These empirical results support the quantitative mechanisms highlighted in the previous subsections. Moreover, they are also consistent with the recent literature on con-

Table A7: Second Stage Regression for Auto Spending at the County Level

	Δ_{06-09} Auto Spending			
FD Measure in 2002:	(DQ30)	(CL80)	(DQ30 and CL80)	(ADQ30)
Δ_{06-09} Home Value	-0.075	-0.090	-0.091	-0.037
	(0.09)	(0.14)	(0.12)	(0.09)
FD	-17.428	-11.568	-14.688	-12.498
	(22.70)	(23.35)	(26.39)	(13.55)
Δ_{06-09} Home Value $\times$ FD	0.830	0.548	0.737	0.420
	(0.54)	(0.51)	(0.59)	(0.34)
Income Volatility	-0.344**	-0.294*	-0.307*	-0.296*
	(0.13)	(0.13)	(0.13)	(0.13)
Unemployment Rate in 2009	-0.121	-0.119	-0.106	-0.047
	(0.30)	(0.31)	(0.31)	(0.34)
% Employment in Tradable Industries	0.114	0.140	0.136	0.121
	(0.11)	(0.11)	(0.11)	(0.12)
Observations	659	659	659	659

Notes: All of the controls from table A2 are included here aggregated to the county level, except for the interactions between the change in home prices and the 2006 levels of income and financial wealth. The regressions are weighted by the number of owner-occupied housing units in the county as of 2006. Standard errors appear in parentheses.

sumption responses to house price shocks as exemplified by Mian, Rao, and Sufi (2013) and Aladangady (2017), among others. However, these results are not intended to establish a causal relationship between financial distress and observed consumption declines. Instead, we are arguing that financial distress is a useful summary statistic capturing a history of high borrowing costs that cannot be accounted for with standard macroeconomic controls. Our model will argue that financial distress is in part driven by an individual's unobservable impatience. The connection between financial distress and vulnerability to shocks shown here corroborates our model's quantitative results, which similarly show that more financially distressed households react more strongly to macroeconomic shocks.

B Recursive formulation of the model

B.1 Nonhomeowner

If the household of type j does not own a house, it must decide whether or not to default on its financial asset/debt holdings a and whether to stay as a renter R or buy a house B. Given these two decisions, we can write the lifetime utility of a household in this situation as:

$$N_{j,n}(a,z,\epsilon) = \max_{I_{rent}\in\{0,1\}} \Bigg\{ I_{rent}R_{j,n}(a,z,\epsilon) + (1-I_{rent})B_{j,n}(a+e_n(z,\epsilon),z) \Bigg\}, \tag{1}$$

where earnings are $e_n(z,\epsilon) = exp(f + l_n + z + \epsilon)$. Here I_{rent} equals 1 when the household chooses to rent, R is the lifetime value of renting, and B is the lifetime value of buying a house. These value functions take the form of:

$$R_{j,n}(a,z,\epsilon) = \max \Bigg\{ R^{P}_{j,n}(a,z,\epsilon), R^{BK}_{j,n}(a,z,\epsilon), R^{DQ}_{j,n}(a,z,\epsilon) \Bigg\}, \tag{2}$$

and

$$B_{j,n}(a,z,\epsilon) = B^{P}_{j,n}(a,z,\epsilon). \tag{3}$$

Notice that households that purchase a house are not allowed to default (in any form) on credit card debt, so the last equality is only for expositional clarity. The superscripts in each value function represent whether the household is, or is not, defaulting on financial assets. We describe these problems next.

Renter and no financial asset default. A household that is a renter and decides *not* to default on financial assets has only to choose the next period's financial assets a':

$$\begin{aligned} R^{P}_{j,n}(a,z,\epsilon) = \max_{a'} \quad & u(c,h_R) + \beta_j \mathbb{E}\Big[N_{j,n-1}(a',z',\epsilon')|z\Big] \qquad (4) \\ s.t. \quad & c + q^{a}_{R,j,n}(a',z)a' = e + a, \\ e \quad = \quad & exp(f + l_n + z + \epsilon). \end{aligned}$$

Here q^{a}_{R} is the price of borrowing financial assets, which depends on the housing state (renter), income states, age, and heterogeneity type j.

Renter and bankruptcy. A household that is a renter and decides to formally default on financial assets a solves the following trivial problem:

$$R^{BK}_{j,n}(a,z,\epsilon) = u(c,h_R) + \beta_j \mathbb{E}\Big[N_{j,n-1}(0,z',\epsilon')|z\Big] \quad (5)$$

$$s.t. \qquad c = e - (\text{filing fee}),$$

$$e = exp(f + l_n + z + \epsilon).$$

Here, the filing fee is the bankruptcy filing fee.

Renter and delinquency. A household that is a renter and decides to skip payments (i.e., become delinquent) on financial assets a solves the following trivial problem:

$$R^{DQ}_{j,n}(a,z,\epsilon) = u(c,h_R) + \beta_j \mathbb{E}\Big[\eta N_{j,n-1}(0,z',\epsilon') + (1-\eta)N_{j,n-1}(a(1+r^R),z',\epsilon')|z\Big] \quad (6)$$

$$s.t. \qquad c = e,$$

$$e = exp(f + l_n + z + \epsilon).$$

Here, η is the probability of discharging delinquent debt and r^R is the roll-over interest rate on delinquent debt.

Homebuyer. A household of type j that is buying a house and has cash in hand a must choose next period's financial assets a', the size of their house h', and the amount to borrow in the mortgage for the house m'.

To simplify the problem later, consider an individual choosing to buy a house of size $h' \in \{h_1 ....., h_m\}$,

$$\hat{B}_{j,n}(a,z;h') = \max_{a',m'} \quad u(c,h') + \beta_j \mathbb{E}\Big[H_{j,n-1}(h',m',a',z',\epsilon')|z\Big] \quad (7)$$

$$s.t. \quad c + q^a_{j,n}(h',m',a',z)a' =$$

$$a + q^m_{j,n}(h',m',a',z)m' - I_{m'>0}\xi_M - (1+\xi_B)ph',$$

$$q^m_{j,n}(h',m',a',z)m' \leq \lambda ph'.$$

Here, q^m is the price of borrowing m' for a house, which depends on house size, income states, and discount factor type j. The other constraints reflect a loan-to-value constraint and that houses must come in discrete sizes. With this notation, the problem of a homebuyer is simply

$$B_{j,n}(a,z) = \max_{h' \in \{h_1,\ldots,h_H\}} \hat{B}_{j,n}(a,z;h'). \tag{8}$$

Notice that in the case of the renter, the cash on hand is simply financial assets plus earnings. Below, we will use the same value function B for individuals in different situations (e.g., moving from one house to another).

B.2 Homeowner

The homeowner's problem is more complex. On the financial asset dimension, homeowners must decide to default or repay their financial assets. On the housing dimension, homeowners can: (i) pay their current mortgage (if any), (ii) refinance their mortgage (or ask for a mortgage if they don't have one), (iii) default on their mortgage, (iv) sell their house and buy another one, or (v) become a renter. The value function H is given by the maximum of:

$$H_{j,n}(h,m,a,z,\epsilon) = \max\left\{ P_{j,n}(\cdot), F_{j,n}(\cdot), D_{j,n}(\cdot), S^B_{j,n}(\cdot), S^R_{j,n}(\cdot) \right\} \tag{9}$$

where:

$$P_{j,n}(h,m,a,z,\epsilon) = \max\left\{ P^P_{j,n}(\cdot), P^{BK}_{j,n}(\cdot), P^{DQ}_{j,n}(\cdot) \right\}, \tag{10}$$

$$F_{j,n}(h,m,a,z,\epsilon) = F^P_{j,n}(\cdot), \tag{11}$$

$$D_{j,n}(h,m,a,z,\epsilon) = \max\left\{ D^P_{j,n}(\cdot), D^{BK}_{j,n}(\cdot), D^{DQ}_{j,n}(\cdot) \right\}, \tag{12}$$

$$S^B_{j,n}(h,m,a,z,\epsilon) = S^{B,P}_n(\cdot), \tag{13}$$

$$S^R_{j,n}(h,m,a,z,\epsilon) = S^{R,P}_n(\cdot). \tag{14}$$

Notice that households that choose to refinance their mortgage cannot default on financial assets in any manner. Additionally, we model agents who elect to sell as having to pay their financial assets.

Mortgage payer and no financial asset default. Households that decide to pay their mortgage and their financial assets have the following problem:

$$P^P_{j,n}(h,m,a,z,\epsilon) \;=\; \max_{a'}\; u(c,h) + \beta_j \mathbb{E}\Big[H_{j,n-1}(h',m(1-\delta),a',z',\epsilon')|z\Big] \tag{15}$$

$$s.t. \qquad c + q^a_{j,n}(h,m(1-\delta),a',z)a' = e + a - m,$$

$$e = exp(f + l_n + z + \epsilon).$$

Mortgage payer and bankruptcy. Households that decide to pay their mortgage but formally default on their financial assets have the following (trivial) problem:

$$P^{BK}_{j,n}(h,b,a,z,\epsilon) \;=\; u(c,h) + \beta_j \mathbb{E}\Big[H_{j,n-1}(h',m(1-\delta),0,z',\epsilon')|z\Big] \tag{16}$$

$$s.t. \qquad c = e - \text{filing fee} - m,$$

$$e = exp(f + l_n + z + \epsilon).$$

Mortgage payer and delinquency. Households that decide to pay their mortgage but choose informal default on their financial assets have the following (trivial) problem:

$$
\begin{aligned}
P_{j,n}^{DQ}(h,m,a,z,\epsilon) \quad &= \quad u(c,h) + \beta_j \mathbb{E}\Big[\eta H_{j,n-1}(h', m(1-\delta), 0, z', \epsilon') \\
&\qquad +(1-\eta) H_{j,n-1}(h', m(1-\delta), a(1+r^R), z', \epsilon')|z\Big]
\end{aligned} \tag{17}
$$

$$
\begin{aligned}
s.t. \qquad & c = e - m, \\
& e = exp(f + l_n + z + \epsilon).
\end{aligned}
$$

Mortgage refinancer. A household that refinances cannot default on financial assets a and must prepay their current mortgage, choose next period's financial assets a', and choose the amount to borrow m' with their new mortgage:

$$
F_{j,n}^{P}(h,m,a,z,\epsilon) = \hat{B}_{j,n}(a + ph(1+\xi_B) - q_n^* m + e_n(z,\epsilon), z; h) \tag{18}
$$

Note that this problem is just a special case of a homebuyer who is "rebuying" their current home of size h but now has cash on hand equal to earnings plus financial assets minus fees from prepaying the previous mortgage m. Also note that $ph(1+\xi_B)$ is simply an adjustment, so the household doesn't actually pay adjustment costs for rebuying their current home.

Mortgage defaulter and no financial asset default. A household that defaults on its mortgage and chooses not to default on its financial assets a immediately becomes a renter and must choose the next period's financial assets a'. Importantly, since we assume the cost of defaulting on a mortgage is a utility cost Φ, we can easily write this problem as the problem of a renter minus the utility cost of mortgage default:

$$
D_{j,n}^{P}(h,m,a,z,\epsilon) = R_{j,n}^{P}(a,z,\epsilon) - \Phi. \tag{19}
$$

Mortgage defaulter and bankruptcy. Using the same trick as above, we can write the problem of a mortgage defaulter who chooses bankruptcy (on financial assets) as the problem of a renter who files for bankruptcy:

$$D^{BK}_{j,n}(h,m,a,z,\epsilon) = R^{BK}_{j,n}(a,z,\epsilon) - \Phi. \tag{20}$$

Mortgage defaulter and delinquency. Lastly, we can write the problem as a mortgage defaulter who chooses delinquency (on financial assets) as the problem of a renter who is delinquent on existing debt:

$$D^{DQ}_{j,n}(h,m,a,z,\epsilon) = R^{DQ}_{j,n}(a,z,\epsilon) - \Phi. \tag{21}$$

Seller to renter. Note that a seller who decides to rent (and not default on financial assets) is simply a renter with financial assets equal to a plus the gains/losses from selling their current house,

$$S^{R,P}_{j,n}(h,m,a,z,\epsilon) = R^{P}_{j,n}(a + ph(1-\xi_S) - q^*_n m, z, \epsilon). \tag{22}$$

Seller to other house. This problem is just a special case of a homebuyer with cash on hand equal to earnings plus current financial assets plus gains/losses from selling the previous house,

$$S^{P,B}_{j,n}(h,m,a,z,\epsilon) = B_{j,n}(a + ph(1-\xi_S) - q^*_n m + e_n(z,\epsilon), z). \tag{23}$$

B.3 Mortgage prices

When a household uses a mortgage that promises to pay m' next period, the amount it borrows is given by $m' q^m_n(h', m', a', z)$, where:

$$q^m_{j,n}(h',m',a',z) \quad = \quad \frac{q^m_{pay,j,n} + q^m_{prepay,j,n} + q^m_{default,j,n}}{1+r}. \tag{24}$$

First, consider the price of payment tomorrow, q_{pay},

$$\begin{aligned} q^m_{pay,j,n}(h',b',a',z) \quad &= \\ &\rho_n \mathbb{E}\Big[\text{mort} \quad \text{pay, no} \quad \text{def} + \text{mort} \quad \text{pay,} \quad \text{BK} + \text{mort} \quad \text{pay,} \quad \text{DQ}\Big| z\Big], \end{aligned} \tag{25}$$

with:

$$\begin{aligned} \text{mort pay, no def} &= \mathbb{I}_{P^P_{j,n-1}(h',m',a',z',\epsilon')}\Big[1 + (1-\delta)q^m_{j,n-1}(h',m'',a^{''},z')\Big], \qquad (26) \\ a^{''} &= \hat{a}^{P,P}_{j,n-1}(h',m',a',z',\epsilon'), \end{aligned}$$

$$\text{mort pay, BK} = \mathbb{I}_{P^{BK}_{j,n-1}(h',m',a',z',\epsilon')}\Big[1 + (1-\delta)q^m_{n-1}(h',m'',0,z')\Big], \qquad (27)$$

and

$$\begin{aligned} \text{mort pay, DQ} &= \mathbb{I}_{P^{DQ}_{j,n-1}(h',m',a',z',\epsilon')}\Big[1 + (1-\delta) \times \qquad (28) \\ &\Big(\eta q^m_{j,n-1}(h',m'',0,z') + (1-\eta)q^m_{j,n-1}(h',m'',a'',z')\Big)\Big], \end{aligned}$$

$$\text{with:} \qquad a'' = (1+r^R)a' \;\text{ and }\; m'' = m'(1-\delta).$$

Here, ρ_n is the age-specific survival probability and $\mathbb{I}$ equals 1 whenever the corresponding value function is the maximum of $P_{j,n-1}$.

Next, consider the price of prepayment tomorrow, q_{prepay}. This occurs when the household chooses to refinance or sell their current house. Importantly, in either case (and regardless of what the household chooses to do immediately after selling their current house) creditors receive value q^*:

$$\begin{aligned} q^m_{prepay,j,n}(h',m',a',z) &= \qquad (29) \\ &\mathbb{E}\Big[\Big(\mathbb{I}_{F_{j,n-1}(h',m',a',z',\epsilon')} \\ &+\mathbb{I}_{S^R_{j,n-1}(h',m',a',z',\epsilon')} + \mathbb{I}_{S^B_{j,n-1}(h',m',a',z',\epsilon')}\Big)q^*_{j,n-1}\Big|z\Big]. \end{aligned}$$

Finally, consider the price of defaulting on the mortgage tomorrow, $q_{default}$. Creditors recover $ph'(1-\bar{\xi}_S)$. So, the price of default is simply:

$$\begin{aligned} q^m_{default,j,n}(h',m',a',z) &= \qquad (30) \\ &\rho_n\mathbb{E}\left[\frac{\Big(\mathbb{I}_{D_{j,n-1}(h',m',a',z',\epsilon')}\Big)ph'(1-\bar{\xi}_S)}{m'}\Bigg|z\right]. \end{aligned}$$

B.4 Bond prices

When a household (that either owns a home or is buying one) issues debt and promises to pay a' next period, the amount it borrows is given by $a' q_n^a(h', m', a', z)$, where:

$$q_{j,n}^a(h', m', a', z) = \frac{q_{pay,j,n}^a + q_{DQ,j,n}^a}{1+r}. \tag{31}$$

First, consider the price of payment tomorrow, q_{pay}^a. This occurs in the following states: homebuyer, no financial asset default; mortgage payer, no financial asset default; mortgage refinancer, no financial asset default; mortgage defaulter, no financial asset default; seller to renter; and seller to buyer. In all of these cases, creditors get paid the same amount per unit of debt issued by the household. Thus,

$$\begin{aligned} q_{pay,j,n}^a(h', m', a', z) = \rho_n \mathbb{E}\Bigg[\quad & \mathbb{I}_{B_{n-1}(a'+e_{n-1}(z',\epsilon'),z',\epsilon')} \\ & + \mathbb{I}_{P_{j,n-1}^P(h',m',a',z',\epsilon')} + \mathbb{I}_{F_{j,n-1}^P(h',m',a',z',\epsilon')} \\ & + \mathbb{I}_{D_{j,n-1}^P(h',m',a',z',\epsilon')} \\ & + \mathbb{I}_{S_{j,n-1}^{R,P}(h',m',a',z',\epsilon')} + \mathbb{I}_{S_{j,n-1}^{B,P}(h',m',a',z',\epsilon')} \Bigg| z \Bigg]. \end{aligned} \tag{32}$$

Next, consider the price given delinquency tomorrow, q_{DQ}^a. This occurs in two states: mortgage payer, delinquency; and mortgage defaulter, delinquency. In both of these cases debt gets rolled over at a rate $(1 + r^R)$ with probability $(1 - \eta)$. However, tomorrow's price of this rolled-over debt varies by state. Thus,

$$\begin{aligned} q_{DQ,j,n}^a(h', m', a', z) = (1-\eta)(1+r^R)\rho_n \mathbb{E}\Bigg[& \mathbb{I}_{D_{j,n-1}^{DQ}(h',m',a',z',\epsilon')} \times q_{R,j,n-1}^a(a'', z') \\ & + \mathbb{I}_{P_{n-1}^{DQ}(h',m',a',z',\epsilon')} \times q_{j,n-1}^a(h', m'', a'', z') \Bigg| z \Bigg] \end{aligned} \tag{33}$$

$$\text{with: } a'' = (1+r^R)a' \qquad \text{and } m'' = m'(1-\delta).$$

Turning to renters, their bond price is simpler. When such a household issues debt and promises to pay a' next period, the amount it borrows is given by $a' q^a_{R,j,n}(a', z)$, where:

$$q^a_{R,j,n}(a', z) = \frac{q^a_{R,pay,j,n} + q^a_{R,DQ,j,n}}{1+r}. \tag{34}$$

Trivially, we have:

$$q^a_{R,DQ,j,n}(a', z) = \rho_n \mathbb{E}\left[\mathbb{I}_{R^P_{j,n-1}(a',z',\epsilon')} \middle| z \right], \tag{35}$$

and

$$q^a_{R,DQ,j,n}(a', z) = (1-\eta)(1+r^R)\rho_n \mathbb{E}\left[\mathbb{I}_{R^{DQ}_{j,n-1}(a',z',\epsilon')} \times q^a_{R,j,n-1}(a'', z') \middle| z \right] \tag{36}$$

$$\text{with:} \quad a'' = (1+r^R)a'.$$

C Restricted models

Table A8 presents the fit of the restricted models discussed in the main text. The columns labeled "Baseline model" present the predictions of the baseline model with both discount factor and rental house size heterogeneity. The column labeled "β-het model" presents the predictions of the restricted model that only allows for discount factor heterogeneity. Lastly, the column labeled "h^R-het model" presents the predictions of the restricted model that only allows for rental house size heterogeneity.

Table A8: Restricted Model Fit by Quintile of FD

	Q1				Q2				Q3				Q4				Q5			
	Data	Baseline model	β-het model	h^R-het model	Data	Baseline model	β-het model	h^R-het model	Data	Baseline model	β-het model	h^R-het model	Data	Baseline model	β-het model	h^R-het model	Data	Baseline model	β-het model	h^R-het model
Savings/Inc	2.44	1.73	1.44	0.34	1.96	1.53	1.29	0.32	1.78	1.35	1.16	0.33	1.57	1.23	1.07	0.31	1.06	1.03	0.91	0.31
Home ownership*	76.30	76.63	87.11	78.91	71.93	71.04	81.90	73.96	68.76	62.34	75.00	65.38	64.25	61.63	69.86	65.86	61.69	52.58	61.49	56.74
Housing leverage	44.11	30.06	33.45	81.13	47.98	36.56	35.03	70.31	44.57	40.89	43.60	76.46	46.04	44.34	43.70	71.42	43.36	44.74	54.53	66.57
Housing debt>0*	49.77	28.75	39.19	70.33	44.67	28.72	40.44	65.60	39.83	25.04	39.30	57.22	36.27	27.43	38.22	57.77	31.84	24.20	36.33	49.17
Housing debt>0* conditional on FD	33.31	32.18	78.08	74.84	30.72	30.75	71.23	70.93	28.37	20.89	63.56	59.52	26.90	27.44	57.90	64.83	25.99	21.37	50.98	53.63
Housing debt/Inc	1.47	0.86	0.80	1.94	1.57	1.03	1.01	1.98	1.57	1.17	1.22	2.19	1.59	1.31	1.37	2.10	1.48	1.52	1.67	2.27
Mortg def rate*	1.52	1.30	1.19	1.56	1.81	1.69	1.44	1.71	2.24	2.20	1.76	2.25	2.58	2.18	2.15	2.12	3.34	2.49	2.47	2.21
DQ rate*	8.98	10.22	9.88	10.83	12.65	13.40	13.21	12.67	15.43	16.22	16.35	12.21	18.28	18.55	18.91	15.06	23.93	22.30	23.00	15.05
BK rate*	0.39	0.59	1.07	0.65	0.55	0.67	1.16	0.66	0.63	0.65	1.11	0.53	0.65	0.69	0.99	0.57	0.64	0.66	0.98	0.53
Persistence of FD:																				
Over 2 yrs	9.20	6.50	4.80	3.56	8.05	5.31	4.09	3.63	6.82	4.87	3.57	4.15	5.89	4.03	3.35	3.53	4.83	3.58	2.87	3.94
Over 4 yrs	6.15	5.30	4.51	2.43	5.36	4.21	3.67	2.40	4.57	3.74	3.09	2.46	3.99	3.17	2.81	2.30	3.20	2.77	2.41	2.42
Over 5 yrs	5.36	5.31	4.53	2.24	4.63	4.18	3.66	2.19	3.98	3.61	3.04	2.16	3.45	3.09	2.77	2.11	2.86	2.68	2.38	2.14
Over 6 yrs	4.86	5.29	4.56	2.17	4.17	4.15	3.67	2.07	3.57	3.58	3.04	1.98	3.20	3.07	2.77	2.00	2.58	2.64	2.37	1.98
Over 8 yrs	3.89	5.13	4.61	1.97	3.56	4.11	3.70	1.94	2.95	3.52	3.03	1.74	2.61	3.01	2.75	1.84	2.19	2.59	2.35	1.75
Over 10 yrs	3.40	4.38	4.52	1.65	3.00	3.68	3.62	1.73	2.66	3.23	2.96	1.47	2.37	2.86	2.73	1.67	2.05	2.47	2.30	1.56
SSE		1.15	5.87	5.66		0.66	3.80	4.49		0.56	2.77	3.94		0.39	2.12	4.27		0.35	1.68	3.10

Notes: $*$ in percent. SSE is the sum of squared errors for each quintile. "Savings/Income" represents mean net financial wealth divided by mean income, and "With housing debt / In FD" is the percent of the population with housing debt, conditional on being in FD.

As can be seen by the bottom row of this table, neither restricted model matches the empirical targets as well as the baseline model. Among restricted models, the β-het model performs the best.

Looking across specific moments shows which heterogeneity is more important to generate certain empirical patterns. For example, the top row reveals that discount factor heterogeneity helps generate a modest decline in savings/income ratios with FD observed in the data. This type of heterogeneity also helps generate the pattern of increasing delinquency rates with FD, as seen in the eighth row. The second to last row shows that discount factor heterogeneity helps generate the relatively elevated persistence of FD at 10 years. Meanwhile, rental house size heterogeneity is helpful in generating the patterns of homeownership and bankruptcy rates by FD that are observed in the data.

For reference, Table A9 presents the parameters used in the restricted models and baseline models. As seen in this table, the columns labeled β-het model highlight the restriction that $h_L^R = h_H^R = \bar{h^R}$ across all quintiles, where $\bar{h^R}$ is equal to the average rental house size across quintiles obtained from the baseline model. Meanwhile, the columns labeled h^R-het model highlight the restriction that $\beta_L = \beta_H = \bar{\beta}$ across all quintiles, where $\bar{\beta}$ is equal to the average discount factor across quintiles obtained from the baseline model.

Table A9: Restricted Model Parameters by Quintile of FD

	Q1			Q2			Q3			Q4			Q5		
Parameter	Baseline model	β-het model	h^R-het model	Baseline model	β-het model	h^R-het model	Baseline model	β-het model	h^R-het model	Baseline model	β-het model	h^R-het model	Baseline model	β-het model	h^R-het model
s_L	0.29	0.29	0.29	0.37	0.37	0.37	0.45	0.45	0.45	0.49	0.49	0.49	0.58	0.58	0.58
h_L^R	4.61	1.54	4.61	3.77	1.54	3.77	3.92	1.54	3.92	2.99	1.54	2.99	2.95	1.54	2.95
h_H^R	0.001	1.54	0.001	0.001	1.54	0.001	0.001	1.54	0.001	0.001	1.54	0.001	0.001	1.54	0.001
β_H	1.00	1.00	0.91	1.00	1.00	0.91	1.00	1.00	0.91	1.00	1.00	0.91	1.00	1.00	0.91
β_L	0.80	0.80	0.91	0.80	0.80	0.91	0.80	0.80	0.91	0.80	0.80	0.91	0.80	0.80	0.91

Notes: This table presents the parameters used for the restricted models and, for comparison purposes, the estimated parameters from the baseline model.

D Alternative model

Here we provide additional details on the alternative model. As alluded to in the main text, the alternative model we use is the same as our baseline model but precludes any delinquency or bankruptcy in the credit market, and thus has no FD. To discipline the amount of borrowing allowed in this economy we introduce an ad hoc borrowing limit $\underline{a}$ (common across quintiles of FD), which we set to the equivalent of 1 times quarterly average labor income, following Kaplan, Moll, and Violante (2018).

We consider two versions of this alternative model. In the first version (what we call the *alternative model* in the main text), we keep the same external and internal parameters the same as in the baseline model. We assume there are five groups of individuals whose preference parameters follow the distribution presented in Table 3. Thus, comparisons between the predictions of this alternative model and the baseline model hold preferences constant and only change the debt market arrangement. However, as seen in Table A10, this model naturally misses many of the empirical targets that the baseline model matches, in part reflecting the importance of FD.

Table A10: Alternative Model Fit by Quintile of FD

	Q1		Q2		Q3		Q4		Q5	
	Data	Alt. Model	Data	Alt. Model	Data	Alt. Model	Data	Alt. Model	Data	Alt. Model
Savings/Inc	2.44	1.66	1.96	1.43	1.78	1.28	1.57	1.09	1.06	0.87
Home ownership*	76.3	70.54	71.93	64.56	68.76	56.43	64.25	55.08	61.69	46.10
Housing leverage	44.11	15.61	47.98	19.00	44.57	18.76	46.04	22.70	43.36	22.21
Housing debt> 0*	49.77	21.82	44.67	21.46	39.83	17.01	36.27	19.72	31.84	15.96
Housing debt/Inc	1.47	0.54	1.57	0.64	1.57	0.64	1.59	0.80	1.48	0.91
Mortg def rate*	1.52	1.08	1.81	0.66	2.24	1.13	2.58	0.76	3.34	0.86

Notes: $*$ in percent. "Savings/Income" represents mean net financial wealth divided by mean income.

Thus, to try to give the alternative model the best chance to match the empirical targets it has in common with the baseline model, we re-estimate it (this is what is referred to as the *re-estimated alternative model* in the main text). Table A11 presents the fit of the re-estimated alternative model, while Table A12 presents the corresponding parameter estimates. For this estimation procedure, we place additional weight on targeting homeownership, since this was one the greatest misses from the un-estimated alternative model.

Looking at Table A11 suggests the alternative model when re-estimated can do a better job of matching the housing-related targets, but struggles to match the dispersion in savings/income ratios across quintiles.

Table A11: Re-estimated Alternative Model Fit by Quintile of FD

	Q1		Q2		Q3		Q4		Q5	
	Data	Alt. Model	Data	Alt. Model	Data	Alt. Model	Data	Alt. Model	Data	Alt. Model
Savings/Inc	2.44	1.31	1.96	1.06	1.78	0.90	1.57	0.89	1.06	0.96
Home ownership*	76.3	75.64	71.93	72.32	68.76	69.24	64.25	64.70	61.69	61.39
Housing leverage	44.11	40.46	47.98	46.73	44.57	50.37	46.04	50.70	43.36	46.64
Housing debt> 0*	49.77	37.18	44.67	40.35	39.83	40.91	36.27	36.31	31.84	30.14
Housing debt/Inc	1.47	1.20	1.57	1.44	1.57	1.57	1.59	1.57	1.48	1.44
Mortg def rate*	1.52	1.24	1.81	1.81	2.24	2.12	2.58	1.76	3.34	1.54
SSE	0.36		0.23		0.27		0.31		0.31	

Notes: $*$ in percent. SSE is the sum of squared errors for each quintile. “Savings/Income” represents mean net financial wealth divided by mean income.

Comparing Table A12 to the baseline model's estimates in Table 3 shows that the alternative model implies a greater fraction of type-L individuals overall. On average, the re-estimated alternative model suggests that 55 percent of the population is of type-L. Meanwhile, the baseline model suggests that 44 percent of the population is of type-L. More salient, though, is the lack of precision in the share estimates (s_L) in this table compared to the baseline model. This highlights the usefulness of FD-related moments in identifying heterogeneity in the population.

Table A12: Alternative Model Parameter Estimates by Quintile of FD

Parameter	Q1	Q2	Q3	Q4	Q5
s_L	0.44	0.55	0.60	0.60	0.56
	(0.21)	(0.31)	(0.41)	(0.32)	(0.22)
h_L^R	3.32	2.37	2.05	2.18	2.24
	(0.17)	(0.10)	(0.05)	(0.04)	(0.29)

Notes: Asymptotic standard errors appear in parentheses.