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	2019-025F
Revision Date	September 2022
Citable Link	https://doi.org/10.20955/wp.2019.025
Suggested Citation	Athreya, K., Mather, R., Mustre-del-Rio, J., Sanchez, J.M., 2022; 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 views expressed in this paper are those of the author(s) and do not necessarily reflect the views of the Federal Reserve System, the Board of Governors, or the regional Federal Reserve Banks. Federal Reserve Bank of St. Louis Working Papers are preliminary materials circulated to stimulate discussion and critical comment.

The Effects of Macroeconomic Shocks: Household Financial Distress Matters*

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

September 14, 2022

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. Thus, quantifying the importance of both dimensions 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.

*We thank seminar participants at University of Rochester, University of New Hampshire, Florida State University, SED Meetings, Bureau of Economic Analysis, Federal Reserve Banks of Atlanta, St. Louis and Kansas City, and Macro System Meetings. The views expressed herein are those of the authors and should not be attributed to the Federal Reserve Banks of Kansas City, Richmond or St. Louis or the Federal Reserve System.

[†]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 the 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 whether and how these two dimensions of heterogeneity matter in shaping consumption responses to aggregate shocks. We make use of a measurement of household financial health we will call financial distress (FD), which 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 Most relevant to the current discussion, we show that during the last two recessions the burden of aggregate shocks was worse among households who were in greater FD prior to each recession.

Gauging the effects of non-uniform shocks and differences in household financial health on consumption 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 formalizes 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.

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

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

³Although the standard model was developed by Huggett (1993) and Aiyagari (1994), to be more precise we are referring here to the more recent quantitative version in Kaplan and Violante (2010) with housing and mortgages.

Next, we construct aggregate shocks to house prices and labor income that mimic the relationship between FD and shocks observed in the last two recessions: the Great Recession (GR) and the COVID-19 pandemic (CV19). Specifically, our calibration of the aggregate shocks replicates the fact that in both episodes the burden of aggregate shocks was worse 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.

In terms of outcomes, we examine both 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 in the cross-section, we additionally consider the *inequality* of consumption responses, and also the impact on the most disadvantaged households, measured by the change in consumption-based *poverty*.⁴

We find that even if house-price or income shocks are uniform across households, heterogeneity in FD still shapes the response of consumption both at the macro- and micro-level. In other words, the uneven burden of aggregate shocks over the past two recessions likely played a secondary role when viewed through the lens of our model.

At the macro-level, accounting for differences in FD across households amplifies the drop in aggregate consumption compared to a model without FD. This is true for both income shocks or house-price shocks. These findings reflect that some households are more responsive to shocks (i.e. have larger MPCs) in a model that accounts for FD compared to one that ignores it.

At the micro-level, the effect of accounting for FD varies with the type of aggregate shock considered. In response to income shocks, consumption inequality and poverty increase by more in a model with FD compared to a model without FD. In response to house-price shocks, both consumption inequality and poverty actually *decrease* in a model with FD, whereas both measures *increase* in a model without FD. This latter result reflects the discipline imposed on the model when estimating it with data on FD. Namely, the data on FD implies that poorer households are more financially constrained and therefore less likely to own houses. As a result, poorer households benefit from

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

declines in house prices as they become more affordable.

Our conclusions on the importance of household FD have practical implications for the literature on “heterogeneity and the macroeconomy.” Our structural estimates may have broader use as they reflect more fundamental—including persistent—differences across agents. Additionally, another key message from our paper is that FD in and of itself is important for understanding the distributional consequences of aggregate shocks because households that differ in FD also differ in their MPCs. In this sense, FD can serve as a “tagging” mechanism to distinguish between high versus low MPC households.

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 model parameterization and estimation, along with the details of calibration of the aggregate shocks. It also provides some validation of the model against external information on the responsiveness of consumption to shocks. With the model developed and vetted, Section 5 contains our main quantitative results, and Section 6 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.

Our examination of 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.⁵ 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

⁵Athreya, Sánchez, Tam, and Young (2015, 2017) and Athreya, Mustre-del Río, and Sánchez (2019) allow for formal and informal default in the financial asset market but have no housing choice.

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 in with exogenous house price shocks follows Corbae and Quintin (2015) and Hatchondo, Martinez, and Sánchez (2015). In that regards, our work is related but different than papers with similar life-cycle model 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 decompositions suggest that much of the effect of FD operates through the structure it imposes on debt holdings across households. Indeed, a model with a no-borrowing constraint (which precludes the discussion of tightening financial constraints) does not generate the responses of consumption to house price shocks that our baseline model does. In this sense, our model is consistent with the view that a significant 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 the FD matters through the ex-ante heterogeneity it encodes and the positive correlation with aggregate shocks to subsequent macroeconomic outcomes.

Next, our conclusion that dispersion in household consumption responses (an endogenous outcome) is mostly due to heterogeneity in MPCs across households (captured by FD) 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 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 a number of 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 find 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.⁶ 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 on the New York Fed Consumer Credit Panel (NY FED-CCP), which contains credit reports for millions of Americans. They are timely because they are updated quarterly and released 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.⁷ 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.

Directly related, Gross and Souleles (2002) use exogenous variation in credit line extensions to gauge the fraction who increase their debt in response (and hence can be viewed as constrained). They find (perhaps unsurprisingly) that those close to their limits increased borrowing by most. A consensus might be that roughly 20% are “constrained” either in terms of excess sensitivity to income or in terms of responses to survey questions.

⁶Any other metrics for FD used within this paper as robustness checks are defined and discussed in appendix section A.3.

⁷Think 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.

Recall that one of our definitions defines FD as close to liquidity constraints: a household is in FD if it has exhausted more than 80% of its credit limit.

Figure 1, taken from Athreya et al. (2019), shows why FD (measured as severe delinquency of 120+ days) is a valuable measure to look at given its relatively high incidence and very high persistence over the life cycle. The dots in this figure show that, on average, roughly 10% of all individuals find themselves in FD, regardless of age. All the other markers reveal that conditional on being in FD today, the likelihood of being in FD in the future is high. For example, the triangles show that conditional on FD today individuals have roughly a 30% probability of being in FD in four years. Given that across all age groups the unconditional probability of FD is roughly 10%, this means that FD today makes it three times more likely to be in FD in four years, compared to unconditionally.

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

Source: Athreya et al. (2019). This figure plots the average probability of being in FD, defined as an individual having a credit card account 120 days or more delinquent at some point during the year.

Aside from being relatively common and very persistent, FD appears to be a valuable measure of vulnerability. For example, Figure 2 reveals an increasing relationship between county-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 county-level).⁸ For ease of exposition, we present the average MPCs for different quintiles of FD, ranging from lowest FD (Q1) to highest FD (Q5).

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

⁸See appendix A.4 for details and robustness of this relationship.

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.

shocks 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 lighter bars in Figure 2 show that this finding holds even when we control for housing leverage. That the result survives the inclusion of housing leverage suggests FD is a broader measure of vulnerability. Intuitively, FD status at any given time encodes information about past debt (non)repayment decisions, something not directly captured by current debt nor leverage. In this sense, FD may help identify households' attitudes toward debt and repayment, which are crucial to determining the consumption response to shocks.

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 at the 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.3.2 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.⁹ 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.¹⁰ As of December 2020, about 25% of individuals in the highest quintile of FD (Q5) reported earnings losses of at least 50%. In contrast,

⁹We are highly appreciative of Alexander Bick and Adam Blandin for sharing their data with us.

¹⁰Graphs with all quintiles appear in the appendix.

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.

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 $\beta_H \geq \beta_L$.¹¹

¹¹Heterogeneity 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).

With respect to 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. In the parametrization section, we will allow for ex-ante differences in the taste for homeownership, which are perfectly correlated with the discount factor heterogeneity. This heterogeneity helps account for observed differences in homeownership across income categories in the United States, and in particular, homeownership differences by FD.

Agents enter each period either as nonhomeowners or homeowners. Rental houses are of size h_R , while owner-occupied houses vary in discrete sizes $h' \in \{h_1, h_2, \dots, h_H\}$. To finance the purchase of nonrental (owner-occupied) houses, agents borrow using mortgages m' . Importantly, borrowing capacity in the mortgage market is endogenously given by a zero-profit condition on lenders due to the limited commitment of agents to repay mortgages.¹²

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.¹³ 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

¹²Housing choices, mortgages, and foreclosures are modeled as in Hatchondo, Martinez, and Sánchez (2015).

¹³Athreya, 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.

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.

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 we use for the value function associated with each choice. For example, the value function for a nonhomeowner with state variable a and y is N . For the sake of brevity, our formal description of this recursive problems is presented in appendix C.

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

$$c + q_{j,n}^a(h_R, 0, a', y)a' = y + a.$$

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 hence on housing, income, and their discount factor type.

Instead, if that renter decides to *formally default* on unsecured debt a , she faces the following trivial budget constraint: $c = y - (\text{filing fee})$, 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 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:

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

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, and the agent's discount factor type j . 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:

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

Figure 6: Decision Tree of a Homeowner

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 following 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:

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

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

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

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 next period's financial assets a' . Thus, the budget constraint she faces is identical to that of a renter who pays her financial assets: $c + q_{j,n}^a(h_R, 0, a', y)a' = y + a$.

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 $c = y - \text{filing fee}$.

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:

$$c + q_{j,n}^a(h_R, 0, a', y)a' = y + a + ph(1 - \xi_S) - q_n^*m.$$

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:

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

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 C.

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 $a'q_{j,n}^a(h', m', a', y)$, where:

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

First, consider the price of a payment tomorrow, $q_{pay,j}^a$. Conditional on being a non-homeowner, this occurs in two scenarios: the agent is a renter with no unsecured debt default or a homebuyer. Conditional on being a homeowner, payment occurs if the homeowner: (i) is a mortgage payer with no unsecured debt default; (ii) is refinancing the mortgage; (iii) is a mortgage defaulter with no unsecured debt default; (iv) is selling the house to become renter; and (v) is selling the house to buy another house. Regardless of homeownership status, in these cases, creditors get paid the same amount per unit of debt issued by the household.

Next, consider the price given delinquency tomorrow, $q_{DQ,j}^a$. Conditional on being a nonhomeowner, this occurs only when renters choose delinquency. Meanwhile, conditional on being a homeowner, 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 bond-pricing formula reveals that bond prices interact with housing status, as the latter affects the likelihood of financial debt payment, default, and delinquency in the future.

4 Model estimation, shock calibration, and model validation

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.¹⁴

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.

Lastly, we evaluate the model’s performance in replicating external information on consumption responses. As seen in Section 4.3, the model’s overall or aggregate MPCs line up with external estimates. Additionally, the model generates a systematic relationship between FD and MPCs that also aligns with the empirical evidence shown earlier in Section 2.1.

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 match key statistics on wealth, home ownership, and FD for each of the five economies.

¹⁴ 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 grouping counties instead of zip codes, and the results are similar.

4.1.1 Assigning first-stage parameters

Table 1 collects the parameters set externally. A period in the model refers to a year. 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%. In addition, we externally calibrate the parameters governing the income process, bankruptcy filing costs, retirement, and mortality. The initial distribution of net financial wealth-to-earnings are 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 $\beta_H=1.00$ and $\beta_L=0.80$, which is within the range of estimates in Athreya, Mustre-del Río, and Sánchez (2019) and also Aguiar, Bils, and Boar (2020). What remains to be determined is the share of people of type- L , s_L , which we pin down below in Section 4.1.2.

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, and θ to 0.11.

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 homeownership (near 100 %) 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. Additionally, 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 $p = 3.3$. 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, 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% of average income, or roughly \$1,000.

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:

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

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

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

We assume $\epsilon_{n,t}^i$ and $e_{n,t}^i$ are normally distributed with variances σ_ϵ^2 and σ_e^2 , respectively. While the income process do not vary across quintiles of FD, the level of income does. 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% larger than the third quintile. Meanwhile, income in the fifth quintile of FD is about 24% 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: $\max\{A_0 + A_1 \exp(z_{W1}^i), A_2\}$. In order to be consistent with U.S. replacement ratios, we calibrate A_0 , A_1 , and A_2 , such that the replacement ratio declines with income, from 69 to 14%, with an average replacement rate of 47%. The age-specific survival probabilities follow Kaplan and Violante (2010).

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_ϵ^2	0.063	Variance of ϵ	Kaplan and Violante (2010)
σ_e^2	0.0166	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 buying a house, households	Gruber and Martin (2003)
ξ_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	Payments 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_H^R	0.001	Size of rental house for patient types	See text
p	3.33	House prices	See text

4.1.2 Estimating the remaining parameters

The remaining parameters to be determined are (i) the share of impatient types in the population s_L ; (ii) the rental house size h_L^R for impatient types; and (iii) the probability of delinquent debt being fully discharged η . We estimate these three parameters so that 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 and shows that the model does well at reproducing the overall patterns. Indeed, the model very closely matches the fact that average delinquency rates rise with each quintile of FD, as do bankruptcy rates. In the model as in the data, the persistence of FD tends to fall over time within a given quintile and, perhaps counter-intuitively, also tends to fall across quintiles as FD increases.

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	1.71	1.96	1.50	1.78	1.36	1.57	1.23	1.06	1.03
Home ownership*	76.3	76.1	71.9	67.8	68.8	62.4	64.2	61.6	61.7	52.8
Housing leverage	44.1	31.1	48.0	37.8	44.6	40.5	46.0	44.6	43.4	44.0
Housing debt > 0*	49.8	28.9	44.7	26.4	39.8	24.9	36.3	27.4	31.8	24.1
Housing debt > 0* conditional on FD	33.3	35.1	30.7	22.8	28.4	21.0	26.9	27.4	26.0	20.9
Housing debt/Inc	1.47	0.87	1.57	1.03	1.57	1.17	1.59	1.31	1.48	1.52
Mortg def rate*	1.52	1.41	1.81	1.63	2.24	2.13	2.58	2.21	3.34	2.49
DQ rate*	8.98	9.64	12.6	13.2	15.4	15.9	18.3	18.5	23.9	22.2
BK rate*	0.39	0.43	0.55	0.58	0.63	0.58	0.65	0.70	0.64	0.65
Persistence of FD:										
Over 2 yrs	9.2	5.15	8.05	5.38	6.82	4.73	5.89	4.07	4.83	3.62
Over 4 yrs	6.15	4.34	5.36	4.16	4.57	3.65	3.99	3.18	3.2	2.80
Over 5 yrs	5.36	4.39	4.63	4.06	3.98	3.54	3.45	3.10	2.86	2.71
Over 6 yrs	4.86	4.48	4.17	4.03	3.57	3.52	3.2	3.07	2.58	2.67
Over 8 yrs	3.89	4.43	3.56	3.95	2.95	3.48	2.61	3.01	2.19	2.62
Over 10 yrs	3.4	3.83	3	3.69	2.66	3.17	2.37	2.84	2.05	2.47
SSE	0.90		0.71		0.57		0.38		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% of the population is

impatient and discounts the future relatively more. In Q5, by contrast, nearly 60% of the population is impatient. Thus, between Q1 and Q5, there is nearly a doubling of this share. The model requires this divergence between Q1 and Q5 to match similarly significant differences between these quintiles in the data. First, the incidence of FD, measured by the DQ rate, is 2.7 times higher in Q5 than in Q1. Second, 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.

Next, the model estimates 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 patient H -types is close to zero, by assumption. The parameter estimates and standard errors in the middle of Table 3 therefore allow us to quickly reject the null of no differences in rental house size between L - and H -types, regardless of quintile of FD. Hence, rental house size heterogeneity is necessary for the model to match the data.

Turning to across quintile differences, the parameter estimates and standard errors also allow us to reject the null of equal rental house sizes for L -type individuals who live in Q1 versus Q5. Interestingly, the model requires a smaller value of h_L^R in Q5 versus Q1. This finding has to do with matching the joint distribution of FD and homeownership. Impatient types tend more often to be in FD, so they must account for a majority of Q5 agents to reproduce the high levels of FD found there in the data. The data also shows that the Q5 homeownership rate is relatively high. However, impatience makes modeled agents less likely to have enough savings to finance home purchases. This discrepancy is resolved in the model by making homeownership comparatively more attractive to L types in Q5 relative to Q1. Hence, the smaller value of h_L^R in Q5 versus Q1.

Lastly, the model also requires significant dispersion across quintiles in the probability of DQ debt being discharged, η . While this probability is 45% in Q1, it is just under 25% in Q5. As previously noted, while L -types will tend more often to be in FD, not all L -types will be in FD. Thus, to incentivize the DQ option in Q1, a higher discharge probability is required. In contrast, since the share of L types is much higher in Q5, the discharge probability is not as high. Beyond affecting the incidence of FD, the discharge

probability also affects its persistence. In an extreme case where very few individuals are in FD (as in Q1), FD would otherwise be very concentrated and persistent within the model. While FD is quite persistent in the data, that persistence quickly declines at longer time horizons. A higher discharge probability helps generate this relatively steep decline within Q1. Again, in contrast, the persistence of FD falls less dramatically in Q5, so the model there requires a lower discharge probability.

Table 3: Parameter Estimates by Quintile of FD

Parameter	Q1	Q2	Q3	Q4	Q5
s_L	0.297 (0.081)	0.385 (0.057)	0.442 (0.054)	0.497 (0.046)	0.575 (0.042)
h_L^R	4.500 (0.016)	4.362 (0.036)	3.943 (0.028)	2.988 (0.035)	2.985 (0.039)
η	0.449 (0.009)	0.294 (0.004)	0.277 (0.003)	0.244 (0.004)	0.244 (0.003)

Notes: Asymptotic standard errors appear in parentheses.

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 data. 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.¹⁵

¹⁵For a rich analysis of the decline in house prices observed during the GFC see Garriga and Hedlund (2017).

Table 4 shows the shocks hitting 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.¹⁶ In terms of labor earnings shocks, we construct the distribution of earnings losses from the survey done by Bick and Blandin (2021).

Table 4: Calibration of House-price and Income Shocks shocks

FD Quintile	Average decline in house prices	Percent of population with earnings loss of:			Average earnings loss
		0%	25%	50%	
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).

These distributions highlight the positive relationship between aggregate shocks and FD. In terms of house prices declines, the Q1 economy (lowest FD) only experiences 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 on severe earnings losses (a 50% decline relative to pre-shock earnings), roughly 14% of households in Q1 receive this type of shock, compared to nearly 22% of households in Q5.

4.3 Model validation

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 externally determined empirics. To do so, we present model-implied MPCs out of house-price and income shocks. 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

¹⁶We obtain very similar results using the average yearly change between 2006 and 2009 as well.

permanent shock. To capture the change over time, we calculate the average annual MPC over three years following the housing shock.¹⁷ We focus our attention on the consumption responses of homeowners. The results are displayed in the first row of Table 5.

Table 5: Model’s MPC Out of House-price and Labor-income Shocks

		(lowest FD)			(highest FD)	
	Aggregate	Q1	Q2	Q3	Q4	Q5
House-price shocks	0.087	0.081	0.081	0.088	0.091	0.095
Labor-income shocks	0.308	0.239	0.287	0.317	0.331	0.385

The first cell of this table shows an aggregate MPC is 8.7 cents per dollar. This number is within the range of estimates in Mian, Rao, and Sufi (2013) in their county-level analysis, who report numbers between 0.054 to 0.119.¹⁸

The others columns of the top row highlight the role of FD in shaping the model implied MPCs out of housing. We observe large dispersion in MPCs across quintiles of FD ranging from roughly 8 cents per dollar (Q1) to 9.5 cents per dollar (Q5). That MPCs out of housing shocks rise with FD is 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 income shocks. The model implies an aggregate MPC of \$0.308 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

¹⁷Calculating MPCs for shorter or longer time horizons does not alter our conclusions.

¹⁸As previously noted, here we focus only on homeowners. If we focus on all the population (as they do), our aggregate MPC is 7 cents per dollar, which is virtually identical to their baseline IV estimate. That homeowners mainly drive the aggregate change in consumption aligns with Aladangady (2017). Our model, however, ascribes a small and subtle effect to renters as well. In the model, renters who will eventually become homeowners experience a small positive income effect from lower house prices 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.

aggregate number, where higher FD is associated with a higher consumption sensitivity per dollar of income lost. In particular, the difference in MPCs between the least and most distressed quintiles (\$0.239 vs. \$0.385) 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. Parker (2017) further argues that “the majority of lack of consumption smoothing is predicted by a simple measure that can be interpreted as impatience.” Perhaps not surprisingly, our structural model generates dispersion in FD via differences in impatience.

5 Quantitative results

In this section we assess the importance of FD in shaping the aggregate and cross-sectional responses of consumption to shocks. To do so, we proceed in two steps. First, we measure how much FD amplifies economic shocks relative to a standard incomplete market model with housing and mortgages. This provides us with a baseline measure of the importance of FD. Second, we disentangle how much of this amplification is due to differences in FD across households versus differences in shocks that are related to FD. We accomplish this by comparing our full model to variants of it which shut-off either differences in FD or differences in shocks.

5.1 The amplification role of FD

As just noted, gauging the amplifying effects of FD can only be done relative to a model without FD. In this subsection we consider a simplified model with no FD and therefore, no differences in shocks correlated with FD. Comparing this simplified model to our full model allows us to quantify the overall importance of FD through all of its channels: because households with different levels of FD have different MPCs (as highlighted in the previous section), and because households in more FD experienced a larger burden of the aggregate shocks during the last two recessions (as documented in the empirical section).

More specifically, our simplified model is a standard heterogeneous agent model with incomplete markets with housing and mortgages. Relative to our full model we remove any notion of FD by imposing a borrowing constraint at zero. Additionally, to stay as close as possible to the standard incomplete markets model we assume agents are ex-ante

identical in their discount factors β . Consequently, there are also no differences in rental house sizes h^R . We calibrate these parameters so that this simplified model matches the savings/income ratio and home ownership rate of the third quintile of FD as reported in Table 2.¹⁹ Finally, because this model has a single type of household, the shocks that it faces are set to match that the aggregate declines in house prices and income from our full model.

Table 6 shows the importance of FD in shaping the responses of aggregate and individual consumption to shocks. The first column of the table collects all of the measures of consumption responsiveness we consider for the full model with FD. The second column of this table displays the corresponding numbers implied by the simplified economy previously described. Finally, the third column presents the difference in responsiveness between the two models, or the amplification that can be attributed to FD through all of its channels.

Table 6: Role of FD for consumption responses to “aggregate” shocks, in percentage deviation from steady state value

	Full model with FD (1)	Simplified model, (2)	Amplification (1)-(2)
Change in consumption p90/p10 ratio			
House-price shocks	-4.45	0.95	-5.40
Labor-income shocks	14.92	2.64	12.28
Change in consumption-based poverty			
House-price shocks	-1.71	2.46	-4.18
Labor-income shocks	17.11	8.29	8.82
Change in aggregate consumption			
House-price shocks	-1.78	-1.08	-0.69
Labor-income shocks	-3.35	-1.47	-1.88

Notes: All values are percentage points of steady-state value. In the housing shock case, these are average changes over three periods following the shock. In the income shock case, the change is measured only in the period of the shock and is calculated over the working-age population since retired agents do not lose any income.

The top panel of this table reveals consequences of each shock on inequality, which we measure with the ratio of the consumption percentiles 90 and 10, referred to as the p90/p10 ratio. A positive (more inequality) change in this p90/p10 ratio after a negative shock indicates that richer households (with ex-ante higher consumption) decrease their consumption by less than poorer households (with ex-ante lower consumption). Table 6

¹⁹Appendix D has more details on the calibration of this model.

shows that in fact this ratio increases in the simplified model after negative shocks to house prices (0.95 percent) or labor income (2.6 percent). Thus, the simplified model indicates that macro shocks generate moderate increases in inequality.

Interestingly, the conclusion is completely different in our full model. Considering the results for house-price shocks, we find that inequality decreases after the shock in our full model with FD. The key difference is that our full model captures the fact that poorer households are less likely to own houses as they are more likely to be debt constrained. For poor households that do not own homes, the decline in house prices is effectively an improvement in housing affordability that allows them to increase non-durable consumption. This mechanism leads the 10th percentile of consumption to increase after a house-price shock, leading to a drop in the $p90/p10$ ratio. In the case of the labor-income shock, the impact on inequality in the full model has the same sign as in the simplified model, but the magnitude is much larger. As will be discussed in the next section, the main reason for this difference is the ability of the full model to capture differences in consumption responses across the consumption distribution.

We also consider the effects that FD has on the lower tail of the consumption distribution in the middle panel of Table 6. Specifically, motivated by the findings in Cutler and Katz (1991), we analyze how FD shapes the response of consumption-based poverty to both house-price and labor-income shocks. We follow Meyer and Sullivan (2019) by targeting a consumption-based poverty threshold of 13% in a steady state, which is the average poverty rate for 2015-2018. Then, we measure how the population share below this threshold changes in response to each shock.

Echoing our findings on consumption inequality, the changes in consumption-based poverty differ dramatically by shock. In the simplified model, poverty increases 2.46 percent after the decline in aggregate house prices. Poverty in our full model, by contrast, *declines* following a house-price decline. Thus, rather than amplifying the house-price shock, modeling FD reverses its effect. While somewhat counter-intuitive, this reversal again reflects that lower house prices increase the affordability of houses for the low-income-high-FD households who do not own houses. Better affordability of houses, in turn, helps alleviate poverty. The critical difference between these models is that in our full model, we capture that low-income households have a high incidence of FD and

low home-ownership rates. Turning to labor-income shocks, consumption-based poverty increases in both models, but more so in our full model with FD. In the simplified model consumption-based poverty increases by roughly 8.3 percent, but poverty increases by more than twice as much in our full model at 17.1 percent.²⁰ Thus, the amplification implied by FD is very significant.

The bottom panel of this table shows that while both models imply significant declines in aggregate consumption following either shock, the decline is always larger in our full model. For reference, the distribution of calibrated shocks in Table 4 implies an aggregate house-price decline of approximately 9 percent and an aggregate earnings losses of 9.6 percent. Focusing on house-price shocks, the simplified model implies an aggregate consumption decline of 1.08 percent, whereas our full model with FD implies a significantly larger drop of 1.78 percent. Overall, our model of FD amplifies the aggregate house-price shock by 0.7 percentage points, or 65 percent more than the simplified model’s response. The implied amplification effect is notably larger when looking at labor-income shocks. In this case, our full model with FD generates an additional decline in consumption of 1.88 percentage points, or 128 percent relative to the simple model without FD.

5.2 Differences in shocks versus differences in households

In the previous section, we demonstrated that including FD amplifies the responses of aggregate consumption, consumption inequality, and poverty to macroeconomic shocks. However, in our full model, FD can amplify responses through a few channels. In this section, we distinguish between mechanisms along which FD influences model outcomes and seek to show which ones are quantitatively more important.

We consider three channels along which FD influences our modeled results. First, our model features unsecured debt and a realistic loan repayment system in which individuals can default formally (through consumer bankruptcy) or informally (through simply not repaying). Modeling FD thus allows households to become debt-burdened and potentially default (formally or informally), and this has consequences for how much they are able

²⁰In other words, consumption-based poverty rises in the full model by 2.22 percentage points from 13 percent to 15.22. This may seem large compared to the true response of poverty to the CV-19 pandemic. However, remember that our model does not include any kind of government programs or intervention. The official income-based poverty rate from the Current Population Annual Social and Economic Supplements (CPS-ASEC) rose by 1ppt to 11.4% in 2020, stopping what had been a 5-year stretch of time in which the annual poverty rate was declining. This suggests that government transfers played a crucial role in shielding consumption from the increase in poverty.

to weather a shock. We refer to this as the *direct* channel of FD.

Second, the data feature a wide dispersion in the prevalence and persistence of default, and so to replicate this pattern, our model also requires for heterogeneous household characteristics. We refer to this as the *indirect* channel because while we use it to match patterns in the data associated with the direct channel, it also has independent implications. For example, people with lower discount factors are less likely to purchase homes regardless of whether or not unsecured debt is available.

Finally, we have noted previously that macroeconomic shocks are experienced differently in different areas, and that for the last two major recessions the hardest-hit areas tended to be the most financially distressed even before the shocks occurred. This is the *correlation* channel.

We isolate these channels through a series of counterfactual economies. Each counterfactual economy is a particular case of our baseline economy with increasingly more simplifying restrictions:

- (a) *Baseline model with uncorrelated shocks:* We compute the size of the shock to house prices and income such that the aggregate decline in house prices and income is the same in our stylized versions of the GR and CV19 recessions. Then, we hit each of the five “regions” with this common shock.
- (b) *“No borrow” model with ex-ante heterogeneous agents and uncorrelated shocks:* We remove FD from the previous model by disallowing *unsecured* borrowing altogether. In practical terms, we impose a zero borrowing constraint.²¹ Although there is no FD (and no correlation with shocks), we still assume there are five different “regions” that differ in preferences as in (a).
- (c) *“Simple Model” with ex-ante identical agents and uncorrelated shocks:* We use the model described in the previous section and referred to as the “simple model.” There are ex-ante identical households, identical shocks, and no unsecured debt.

We attribute the difference between the baseline model and (a) to the *correlation channel*,

²¹We also did this exercise in an alternative model having unsecured debt but disallowing default. The predictions regarding the consumption response to shocks are similar to the model with no borrowing presented here.

the difference between (a) and (b) to the *direct channel*, and the difference between (b) and (c) to the *indirect channel*.

Table 7 summarizes the results of this decomposition and reveals three broad conclusions. First, when considering house-price shocks, across all measures the *direct* channel of FD is key. Second, when considering labor-income shocks, the *indirect* channel of FD is the most important. Finally, the *correlation* channel of FD is of less importance and this holds true across shocks and measures considered. This last observation is the basis of our conclusion that differences across households appear to be more important than differences in shocks.

Looking at the house-price shock rows, it is immediately apparent that amplification of these types of shocks is disproportionately due to the *direct* effect of FD. When looking at the change in the p90/p10 ratio of consumption, the direct channel accounts for more than 100 percent of all of the amplification described in the previous section. In the case of consumption-based poverty, the direct channel still accounts for over 80 percent the amplification of FD, with the indirect channel contributing 18 percent. In terms of the aggregate change in consumption, the direct channel contributes nearly 90 percent of the total amplification effect of FD. Lastly, across each of these rows, the correlation channel's contribution is always negative and comparatively small.

The top panel of Figure 7, which plots changes in consumption percentiles for each model following a house-price shock, helps illustrate the p90/p10 ratio results graphically. Looking at the size and direction of the bars representing the changes in the 10th and 90th percentile, reveals that movements in the 10th percentile are key for delivering the changes in the p90/p10 ratio that we previously described.

Looking at the first two bars (baseline model with correlated shocks and uncorrelated shocks, respectively) of the 10th percentile category reiterates that the correlation channel is relatively unimportant as the implied change in this percentile between our baseline model with correlated shocks and the same model with uncorrelated (or uniform) shocks is essentially the same. In contrast, once we remove FD (the third bar) the change in the 10th percentile reverses sign. This suggests the direct channel of FD is key in generating the decline in consumption inequality documented in the previous section. That some of the sign reversal vanishes once we remove any preference heterogeneity, as in the fourth

bar, highlights the smaller importance of the indirect channel of FD for the amplification of house-price shocks.

These same movements in the 10th percentile of consumption provide some intuition for the results on consumption-based poverty previous discussed. Recall, the poverty threshold is such that 13 percent of the population is in poverty in steady state. Hence, movements in the 10th percentile of consumption are informative. Indeed, given that the 10th percentile barely changes whether shocks are correlated or not is consistent with the correlation channel of FD being unimportant for the amplification of house-price shocks into poverty. By contrast, that the 10th percentile increases in the second bar, but falls in the third bar is consistent with the direct channel of FD driving the decrease in poverty following a house-price shock as documented in the previous section.

Finally, the top panel of Figure 7 also sheds light on the how FD amplifies the effects of house-price shocks on aggregate consumption. Across all models, the top 25 of the consumption distribution accounts for a disproportionate share of total consumption (roughly 60 percent), whereas the bottom 25 accounts for a much smaller share (less than 10 percent). Thus, in order to understand the response of aggregate consumption it suffices to understand how spending patterns of the top 25 change.

As can be seen from the top panel of this figure, the correlation channel of FD matters little in aggregate as the first two bars (baseline model with correlated or uncorrelated shocks) in the 75th, 90th, and 95th quintiles are all very similar. However, once we remove FD, as in the third or fourth bars, we see larger differences in the response of the 90th and 95th percentiles. In the case of the 95th percentile, the simple model with no FD, no heterogeneity, and uncorrelated shocks implies no change. In contrast, even with model with no FD but heterogeneity implies a larger contraction. Thus, the direct channel of FD is critical in amplifying the drop in aggregate consumption.

Turning to the labor-income shock rows, it is also apparent that amplification of these types of shocks is disproportionately due to the *indirect* effect of FD. When looking at the change in the p90/p10 ratio of consumption, the indirect channel accounts for more than 100 percent of all of the amplification described in the previous section. In the case of consumption-based poverty, the indirect channel accounts for over 90 percent the amplification of FD. In terms of the aggregate change in consumption, the indirect channel

Figure 7: Response of consumption percentiles to shocks, alternative model

Table 7: Contribution of FD’s channels for the main results, as percent of the total amplification

	Direct	Indirect	Correlation
Change in consumption p90/p10 ratio			
House-price shock	108.83	-7.55	-1.28
Labor-income shock	-24.55	112.05	12.50
Change in consumption-based poverty			
House-price shock	83.21	18.33	-1.54
Labor-income shock	0.01	91.96	8.04
Change in aggregate consumption			
House-price shock	88.91	19.57	-8.47
Labor-income shock	14.17	81.13	4.70

Notes: Each number is the ratio of (i) the extra effect because of the channel and (ii) the total amplification from Table 6.

contributes over 80 percent of the total amplification effect of FD. Lastly, the correlation channel’s contribution is positive, but still not the dominant source of amplification.

The bottom panel of Figure 7 again helps illustrate many of these points. Like in the case of house-price shocks, this figure shows the 90th percentile changes much less compared to the 10th percentile, so movements in the latter drive movements in the p90/p10 ratio following a labor-income shock. Comparing the first two bars of the 10th percentile shows the correlation channel’s positive contribution to amplification. Indeed, removing correlated shocks exacerbates the decrease in the 10th percentile. Removing FD, as in the third bar, barely changes the drop in the bar. Rather, removing preference heterogeneity, as in the fourth bar, reduces the drop in the 10th percentile significantly. Hence, the indirect channel of FD is critical in amplifying the effects of labor-income shocks into consumption.

Once again, these same movements in the 10th percentile of consumption provide some intuition for the results on consumption-based poverty when a labor-income shock hits. Recall, poverty rises when a labor-income shock arrives and this rise is significantly larger in our baseline model with FD compared to a simple model without FD, correlated shocks, or preference heterogeneity. Consistently, the drop in the 10th percentile is largest in models with either FD or the heterogeneity it encodes (the second and third bars) and much smaller in the simple model (fourth bar). Importantly, ignoring FD but modeling the heterogeneity it encodes (the third bar) is sufficient to replicate the drop in the 10th percentile the baseline model generates, suggesting the indirect channel of FD drives the

amplification of labor-income shocks into consumption-based poverty.

Finally, the bottom panel of Figure 7 also explains why the drop in aggregate consumption is larger in our baseline model with FD than in our simpler model. Indeed, the simplest model with no FD (the fourth bar) generates the smallest drops in consumption across all percentiles of the consumption distribution. Critically, though the indirect channel of FD is driving these results. Adding back in the heterogeneity that FD encodes, as in the third bar, delivers the majority of the drop in consumption across percentiles that the baseline model with correlated shocks generates.

6 Concluding remarks

Our paper aims to understand how household financial distress (FD) shapes individual and aggregate consumption dynamics after aggregate shocks. FD 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, over the past two recessions exposure to aggregate shocks was correlated with prior FD.

We find that FD plays an essential role in understanding how shocks propagate to consumption because it implies significant heterogeneity across households, which itself leads to significant heterogeneity in consumption responses. That prior FD was correlated with aggregate shock exposure matters less for understanding how consumption responds. More specifically, when focusing on house-price shocks, FD matters because of the discipline it places on debt arrangements. In particular, models with FD imply that poorer households are less likely to own homes, consistent with the data. When focusing on labor-income shocks, FD also matters, but not because of the discipline it places on debt arrangements. Rather, matching the persistence of FD requires a significant degree of ex-ante preference heterogeneity across households. This heterogeneity implies that some households are more responsive to shocks (have higher MPCs) compared to a model that ignores FD and its persistence.

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 debt forbearance. Additionally, our model can be used to understand survey results in Coibion, Gorodnichenko,

and Weber (2020) 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., 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.
- Berger, D. and Vavra, J. Shocks versus Responsiveness: What Drives Time-Varying Dispersion? *Journal of Political Economy*, 127(5):2104–2042, 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.
- 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.
- Cutler, D. M. and Katz, L. F. Macroeconomic Performance and the Disadvantaged. *Brookings Papers on Economic Activity*, 22(2):1–74, 1991.
- 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.
- Gross, D. B. and Souleles, N. S. Do liquidity constraints and interest rates matter for consumer behavior? evidence from credit card data. *Quarterly Journal of Economics*, 117(1):149–185, February 2002.
- Gruber, J. and Martin, R. Precautionary savings and the wealth distribution with illiquid durables. 2003.
- 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.
- 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., 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.
- 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.

Online Appendix

A	Empirical analysis	43
A.1	A geographically representative sample	43
A.2	Constructing measures of wealth and consumption	44
A.3	Financial distress	45
A.4	Regressions	46
B	Income Inequality	51
C	Recursive formulation of the model	52
C.1	Nonhomeowner	52
C.2	Homeowner	57
C.3	Mortgage prices	60
C.4	Bond prices	62
D	Calibration of simplified model	63
E	More on alternative models	63

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 size²² 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 pulling a 50 percent sample of Equifax records for the largest zip codes.²³ In order to remain in our sample for a given quarter, individuals must be between 25 and 65 years old, inclusive.²⁴

²²Specifically by using the “number of returns” field provided by the IRS SOI.

²³Zip-code level data on CL80 and DQ30 are available at this link for the years 2006 and 2018.

²⁴Age is calculated using an individual’s recorded birth year, and so any records not including a birth year are also excluded.

Then, we 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.²⁵ 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 debt²⁶ 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.²⁷ FW at the county level is simply calculated as the sum of FW in its component zip codes.²⁸ D is calculated in a similar fashion 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.²⁹ Next, we assign each geography a portion of the total debt from the SCF equal to that fraction.

²⁵To 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.

²⁶Here we include mortgages, the home equity installment balance, and the home equity revolving balance.

²⁷Mian, 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.

²⁸To 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.

²⁹Because our method of pulling Equifax data intentionally over sampled 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 also use a measurement of consumption used in the literature, the new of new cars registrations. In particular, we use data from R.L. Polk by IHS Markit to find the quantity of new automobiles registered in 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 follow after 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

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.

A.3.1 The persistence of “pre-existing” regional FD

FD defined in this way is highly persistent over time at an individual level, as shown in Athreya, Mustre-del Río, and Sánchez (2019). Thinking of a zip code as a collection of individuals, it follows that there should be some slight persistence in FD characteristics at a community level as well, limited by the way that individuals sometimes move. In fact, however, FD at the zip-code level is much more persistent than would be expected if individual-level persistence were the only factor at play.

Given from the 2007 ACS data that the average person will move about 12 times in their lifetime, and assuming that those moves are distributed randomly over 80 years, a back of the envelope calculation suggests that this average person would move 2.6 times in the years 2001-2017. If there were no tendency for people to sort themselves into zip codes with similar FD patterns or be somehow influenced into FD by their surrounding community, we would therefore expect that a zip code’s FD in 2000 would carry little

predictive power for its status in 2017. Conditional on a zip code having been in the worst quintile of FD in 2000, however, there is a 55 percent chance that it was still in the worst quintile 18 years later. This is over twice as likely as random chance would predict. In addition, zip codes that did leave the worst quintile did not move far: 24 percent had moved to quintile 4 by 2018, and only 4 percent had moved to the least-distressed quintile.

The persistence of regional FD helps us to disentangle the underlying pre-existing conditions of FD at the onset of an economic shock from an FD response endogenously made due to the shock. For each shock we consider, distinguishing zip codes that temporarily entered FD in this way from those that were already in FD requires measuring FD somehow separately from this endogenous response. Because FD is so persistent, this can be done by measuring it within each zip code before the shock occurred. We specifically use FD measurements taken in 2002 for the housing shock modelling the Great Recession and measurements taken in 2018 for the income shock modelling the COVID-19 pandemic.

A.3.2 Robustness of correlation between FD and the house-price Shock

First, we show that the correlation we established in Figure 3 in the main text holds if we replace home values by housing wealth shocks as in Mian, Rao, and Sufi (2013). Figure A1 documents the 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 regional persistence of FD discussed in appendix Section A.3.1, 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.

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

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 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 the year of pre-existing conditions

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.

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 limit³⁰ and then 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

³⁰To 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 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.

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

FD Measurement taken in 2002:	(DQ30)	Δ_{06-09} Auto Spending (CL80)	(CL80 and DQ30)	(ADQ30)
Δ_{06-09} Home Value	-0.005 (0.00)	-0.008 (0.00)	-0.009* (0.00)	-0.006 (0.00)
FD	-5.283*** (1.15)	-5.203*** (1.02)	-5.525*** (1.19)	-3.670*** (0.74)
Δ_{06-09} Home Value $\times$ FD	0.099*** (0.02)	0.070*** (0.02)	0.097*** (0.02)	0.070*** (0.01)
Observations	14136	14136	14136	14136

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 percent of households that owned homes in 2006 and include 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 does not appear to be any clear contemporaneous relationship between FD and housing leverage in 2002. Considering the 2006 housing leverage ratio against FD in 2002, there appears to be if anything a negative relationship between the two; i.e., regions with more financial distress tend to have lower leverage.

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 tables A4 and A5, where we present the first and second stages of the regression as we do in Table A2 but instead at the county level.

These empirical results support the quantitative mechanisms highlighted in the previ-

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

FD Measurement taken in 2002:	Δ_{06-09} Auto Spending			
	(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 $\times$ 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 $\times$ Housing Leverage Ratio ₀₆	0.018*	0.014	0.016*	0.019*
	(0.01)	(0.01)	(0.01)	(0.01)
Housing Leverage Ratio, 2006 $\times$ 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.

ous subsections. Moreover, they are also consistent with the recent literature on consumption 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

Table A4: First-Stage Regression, County-level data

FD Measurement taken in 2002:	(DQ30)	Δ_{06-09} Home Value		
		(CL80)	(CL80 and DQ30)	(ADQ30)
Saiz Elasticity	19.586*** (1.80)	19.900*** (1.80)	19.705*** (1.80)	18.762*** (1.83)
FD	109.420* (52.50)	43.793 (51.03)	91.864 (58.10)	91.533** (33.22)
Observations	670	670	670	670
F	31.97	31.47	31.47	32.45

Table A5: Second Stage regression, County-Level

FD Measurement taken in 2002:	(DQ30)	Δ_{06-09} Auto Spending		
		(CL80)	(CL80 and DQ30)	(ADQ30)
Δ_{06-09} Home Value	-0.273* (0.11)	-0.442** (0.16)	-0.367* (0.15)	-0.288* (0.12)
FD	-27.774 (19.83)	-23.662 (19.55)	-29.674 (22.50)	-17.526 (12.03)
Δ_{06-09} Home Value $\times$ FD	1.260* (0.53)	1.304* (0.52)	1.407* (0.59)	0.887* (0.35)
Observations	670	670	670	670

Notes: Regressions are weighted by the number of owner-occupied housing units in each County in 2006. Additional controls not shown here include interactions between the levels and changes in housing wealth, income, and financial wealth.

a causal relationship between financial distress and observed consumption declines. Our model suggests that financial distress is a useful summary statistic capturing a history of high borrowing costs induced in part by individual impatience, which is difficult to observe directly. Rather, these results corroborate our model’s quantitative implications.

B Income Inequality

Similar empirical results can be obtained looking at income inequality as the factor differentiating regions instead of FD.

First, we mentioned in the paper that there is a dispersion in FD across zip codes. Clearly, there is income inequality across zip codes, as shown in Figure A5.

Second, we showed that there is a positive correlation between “aggregate shocks” and the incidence of FD. Figure A6 shows this relationship again but now replacing FD with per household income. The two top figures show that the fact is also true for income per household. The bottom plots what it may be obvious: FD and income per household are very correlated.

Figure A5: Lorenz Curve of Income Across Zip Codes

Our analysis in the main body of the paper continues by making 5 quintiles of FD and calibrating an economy for each quintile. In Table A6 below we show that the quintiles would look similar if we made them according to income per household. To compare with FD quintiles it is useful to compare DQ30. Notice that when the quintiles are made according to income per households, we find that DQ30 decreases from 20.3 percent to 10.8 percent. Obviously, the difference is larger when we made the quintiles according to DQ30 but the difference is not that significant. In that case DQ30 in Q1 is 8.6 and in Q2 is 23.5.

We could continue the analysis using quintiles of income to calibrate five economies and compute the quantitative results. However, we think that the numbers presented here are enough to conclude that we would obtain very similar results, perhaps slightly weaker because the difference in FD are a bit smaller.

C Recursive formulation of the model

C.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:

Figure A6: Macro risks and income inequality

Change in house prices 2006-2012 and per household income 2002

Workers Leisure and Hospitality and per household income 2018

Relationship between financial distress and income, 2002

Table A6: Quintiles Of Income

	(1)				
	1	2	3	4	5
Income Per Household \$000	31.69	41.83	51.54	65.40	106.2
Net Wealth Per Hou. \$000 SCF	82.86	116.8	144.3	191.7	415.6
Fin. Wealth Per Hou. \$000 SCF	83.32	113.7	136.8	177.0	373.6
Net Fin. Wealth Per Hou. \$000 SCF	48.22	67.07	78.61	101.7	266.5
Median Home Value \$000	115.6	143.4	172.1	218.0	328.2
Less Than HS	25.79	19.57	15.93	11.23	7.113
HS	29.61	29.53	27.65	24.24	17.50
College	44.61	50.90	56.42	64.52	75.39
Age	42.75	43.06	43.24	43.47	44.12
Percent that Own a Home	56.71	64.23	69.31	74.59	78.08
Percent with Mortgage/HELOC Debt	29.26	36.13	40.68	46.11	50.22
Housing Debt per Home Owner \$000	53.81	66.74	80.40	101.2	145.0
CC Debt Per Household \$000	3.376	3.962	4.386	4.934	5.770
Housing Leverage	46.60	44.47	44.46	46.73	43.82
Percent with FD who have Mortgage/HELOC debt	23.52	26.84	29.28	31.66	33.99
Percent with Mortgage/HELOC Debt in FD	16.37	13.48	11.72	9.314	7.171
Foreclosure Rate	2.999	2.498	2.288	1.939	1.759
Bankruptcy Rate	0.662	0.694	0.638	0.515	0.353
DQ30	20.27	18.15	16.32	13.71	10.80

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

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:

$$R_{j,n}(a, z, \epsilon) = \max \left\{ R_{j,n}^P(a, z, \epsilon), R_{j,n}^{BK}(a, z, \epsilon), R_{j,n}^{DQ}(a, z, \epsilon) \right\}, \quad (2)$$

and

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

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 next period's financial assets a' :

$$R_{j,n}^P(a, z, \epsilon) = \max_{a'} u(c, h_R) + \beta_j \mathbb{E} \left[N_{j,n-1}(a', z', \epsilon') | z \right] \quad (4)$$

$$s.t. \quad c + q_{j,n}^a(h_R, 0, a', z) a' = e + a,$$

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

Here q^a is the price of borrowing financial assets, which depends on housing, income states, and discount factor 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:

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

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

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

Here, 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:

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

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

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

Here, γ 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 a individual choosing to buy a house of size $h' \in \{h_1, \dots, h_m\}$,

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

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

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

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

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

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

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).

C.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:

$$H_{j,n}(h, m, a, z, \epsilon) = \max \left\{ P_{j,n}(\cdot), F_{j,n}(\cdot), D_{j,n}(\cdot), S_{j,n}^B(\cdot), S_{j,n}^R(\cdot) \right\} \quad (9)$$

where:

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

$$F_{j,n}(h, m, a, z, \epsilon) = F_{j,n}^P(\cdot), \quad (11)$$

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

$$S_{j,n}^B(h, m, a, z, \epsilon) = S_n^{B,P}(\cdot), \quad (13)$$

$$S_{j,n}^R(h, m, a, z, \epsilon) = S_n^{R,P}(\cdot). \quad (14)$$

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 also 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:

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

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

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

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

$$P_{j,n}^{BK}(h, b, a, z, \epsilon) = u(c, h) + \beta_j \mathbb{E} \left[H_{j,n-1}(h', m(1 - \delta), 0, z', \epsilon') | z \right] \quad (16)$$

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

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

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

$$P_{j,n}^{DQ}(h, m, a, z, \epsilon) = u(c, h) + \beta_j \mathbb{E} \left[\gamma H_{j,n-1}(h', m(1 - \delta), 0, z', \epsilon') \right. \\ \left. + (1 - \gamma) H_{j,n-1}(h', m(1 - \delta), a(1 + r^R), z', \epsilon') | z \right] \quad (17)$$

$$s.t. \quad c = e - m,$$

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

Mortgage refiner. 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:

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

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 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:

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

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

$$D_{j,n}^{BK}(h, m, a, z, \epsilon) = R_{j,n}^{BK}(a, z, \epsilon) - \Phi. \quad (20)$$

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:

$$D_{j,n}^{DQ}(h, m, a, z, \epsilon) = R_{j,n}^{DQ}(a, z, \epsilon) - \Phi. \quad (21)$$

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,

$$S_{j,n}^{R,P}(h, m, a, z, \epsilon) = R_{j,n}^P(a + ph(1 - \xi_S) - q_n^*m, z, \epsilon). \quad (22)$$

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,

$$S_{j,n}^{P,B}(h, m, a, z, \epsilon) = B_{j,n}(a + ph(1 - \xi_S) - q_n^*m + e_n(z, \epsilon), z). \quad (23)$$

C.3 Mortgage prices

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

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

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

$$q_{pay,j,n}^m(h', b', a', z) = \rho_n \mathbb{E} \left[\text{mort pay, no def} + \text{mort pay, BK} + \text{mort pay, DQ} \middle| z \right], \quad (25)$$

with:

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

$$\text{mort pay, BK} = \mathbb{I}_{P_{j,n-1}^{BK}}(h', m', a', z', \epsilon') \left[1 + (1 - \delta) q_{j,n-1}^m(h', m'', 0, z') \right], \quad (27)$$

and

$$\begin{aligned} \text{mort pay, DQ} &= \mathbb{I}_{P_{j,n-1}^{DQ}}(h', m', a', z', \epsilon') \left[1 + (1 - \delta) \times \right. \\ &\quad \left. \left(\gamma q_{j,n-1}^m(h', m'', 0, z') + (1 - \gamma) q_{j,n-1}^m(h', m'', a'', z') \right) \right], \end{aligned} \quad (28)$$

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

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^* :

$$\begin{aligned} q_{prepay,j,n}^m(h', m', a', z) &= \\ &\mathbb{E} \left[\left(\mathbb{I}_{F_{j,n-1}}(h', m', a', z', \epsilon') \right. \right. \\ &\quad \left. \left. + \mathbb{I}_{S_{j,n-1}^R}(h', m', a', z', \epsilon') + \mathbb{I}_{S_{j,n-1}^B}(h', m', a', z', \epsilon') \right) q_{j,n-1}^* \middle| z \right]. \end{aligned} \quad (29)$$

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:

$$q_{default,j,n}^m(h', m', a', z) = \rho_n \mathbb{E} \left[\frac{\left(\mathbb{I}_{D_{j,n-1}}(h', m', a', z', \epsilon') \right) p h' (1 - \bar{\xi}_S)}{m'} \middle| z \right]. \quad (30)$$

C.4 Bond prices

When a household issues debt and promises to pay a' next period, the amount it borrows is given by $a' q_n^a(h', b', a', z)$, where:

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

First, consider the price of payment tomorrow, q_{pay}^a . This occurs in the following states: renter, no financial asset default; 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,

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

Notice that the first two terms of the expectation can only occur if $h' = h_R$, whereas the latter five only occur if $h' > h_R$. Additionally, the first default term is unnecessary since mortgage default never occurs without the depreciation shock when house prices are constant.

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

$$q_{DQ,j,n}^a(h', m', a', z) = (1 - \gamma)(1 + r^R)\rho_n \mathbb{E} \left[\begin{aligned} &\mathbb{I}_{R_{j,n-1}^{DQ}(a', z', \epsilon')} \times q_{j,n-1}^a(h_R, 0, a'', z') \\ &+ \mathbb{I}_{D_{j,n-1}^{DQ}(h', m', a', z', \epsilon')} \times q_{j,n-1}^a(h_R, 0, a'', z') \\ &+ \mathbb{I}_{P_{n-1}^{DQ}(h', b', a', z', \epsilon')} \times q_{j,n-1}^a(h', m'', a'', z') | z \end{aligned} \right] \quad (33)$$

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

Notice here too that the first term can only occur if $h' = h_R$, whereas the latter two only occur if $h' > h_R$.

D Calibration of simplified model

Table A7 presents the fit of the simple model to moments we target. Because we have only two parameters to pin down (the discount factor and rental house size), we target the savings/income ratio and home ownership rate, both of the third quintile of FD. Table A8 presents the implied parameter values from this calibration procedure.

Table A7: Fit of simplified model

	Data	Model
Savings/Inc	1.78	1.84
Home ownership*	68.8	73.4

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

Table A8: Parameter Estimates of simplified model

Parameter	Value
β	0.99
h^R	2.00

E More on alternative models

Table A9 presents the key statistics for each of the alternative models. These numbers are used to compute the decomposition of the amplification in direct, indirect, and correlation channels.

We also provide some additional details and results on the *“no borrow” model with ex-ante heterogeneous agents and uncorrelated shocks* and the *“simple model” with ex-*

Table A9: Aggregate Consumption Changes By Model and Shock

	Baseline corr. shocks	Baseline unc. shocks	No-borrow unc. shocks	Simplified unc. shocks
Change in consumption p90/p10 ratio				
House-price shocks	-3.50	-3.22	0.46	1.30
Labor-income shocks	14.92	13.39	16.40	2.64
Change in consumption-based poverty				
House-price shocks	-2.16	-2.71	-4.31	3.59
Labor-income shocks	17.11	16.40	16.40	8.29
Change in aggregate consumption				
House-price shocks	-1.78	-1.83	-1.22	-1.08
Labor-income shocks	-3.35	-3.26	-2.99	-1.47

Notes: All values are measured as percentage changes relative to the old steady-state. In the housing shock case, these are average changes over three periods following the shock. In the income shock case, the change is measured only in the period of the shock and is calculated over the working-age population since retired agents do not lose any income.

ante identical agents and uncorrelated shocks. As noted in the main text, both of these models impose a zero borrowing constraint, which effectively disallows the possibility of FD as we define it.

In the case of the “no borrow” model with *ex-ante heterogeneous agents and uncorrelated shocks*, we assume five different “regions” exist each with parameters following Table 3, but implicitly with the restriction that $\eta = 0.0$ because there is no FD. Since we assume equal shocks, each region is subject to the same distribution of shocks in the corresponding experiment.