



Working Paper Series

Household Financial Distress and the Burden of 'Aggregate' Shocks

WP 20-12

Kartik Athreya
Federal Reserve Bank of Richmond

Ryan Mather
Federal Reserve Bank of St. Louis

Jose Mustre-del-Rio
Federal Reserve Bank of Kansas City

Juan M. Sanchez
Federal Reserve Bank of St. Louis

This paper can be downloaded without charge
from: <http://www.richmondfed.org/publications/>



Richmond • Baltimore • Charlotte

Household Financial Distress and the Burden of “Aggregate” Shocks^{*}

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

September 15, 2020

Abstract

The goal of this paper is to show that household-level financial distress (FD) varies greatly, meaning there is unequal exposure to macroeconomic risk, and that FD can increase macroeconomic vulnerability. To do this, we first establish three facts: (i) regions in the U.S. vary significantly in their “FD-intensity,” measured either by how much additional credit households therein can access, or in how delinquent they typically are on debts, (ii) shocks that are typically viewed as “aggregate” in nature hit geographic areas quite differently, and (iii) FD is an economic “pre-existing condition”: the share of an aggregate shock borne by a region is positively correlated with the level of FD present at the time of the shock. Using an empirically disciplined and institutionally rich model of consumer debt and default, we show that in the shocks dealt by the Great Recession and in the initial months in the COVID-19 pandemic, FD mattered. Our model implies that the uneven distribution of FD creates widely varying consumption responses to shocks. This is true even when subjecting regions (with differing levels of FD) to the same shocks, which highlights the importance of FD independently of its correlation with shocks.

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

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

^{*}We thank seminar participants at the St. Louis Fed and Kansas City Fed, 2018 A. Stockman Conference, SED conference, and Atlanta Fed “Joint Central Bankers Conference.” 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 Richmond; e-mail: kartik.athreya@rich.frb.org

[‡]Federal Reserve Bank of St. Louis; e-mail: ryan.mather@stls.frb.org.

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

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

1 Introduction

The primary goal of this paper is use data and quantitative theory to measure the presence—and inequality of—financial distress (FD) across households and assess the role it plays in how households respond to macroeconomic shocks. We will also, in turn, assess how the aggregate state of financial distress in the economy matters for macroeconomic outcomes.

Our paper has both an empirical and a structural component. Starting with the data, we establish three main facts regarding financial distress in U.S. macroeconomic data. First, using proprietary data, we show that the prevalence of household balance sheets with FD—as measured either by proximity to exhaustion of available credit commitments from lenders, or as having delinquent debt—varies across regions. Second, we establish from the data that the “localized” impact of an event associated with a major aggregate disruption, such as the Great Recession or the COVID-19-induced downturn, also varies greatly. Third, we show that the size of these localized shocks is positively correlated with the level of FD present at the time of the shock. In other words, the most financially distressed households are also, apparently, the least fortunate in terms of exposure to recessions.

To understand what these facts imply for households and the macroeconomy, we develop a sophisticated model of household consumption, debt, and default decisions. We use the model, once disciplined by the facts, to measure the impact of an empirically positive correlation between the aggregate shocks hitting a region and that region’s level of FD at the time of the shock—hereafter noted as its *prior* FD. Specifically, we examine how local responses and macroeconomic conditions are affected by the demonstrated positive correlation of a zip code’s prior FD with (i) asset price declines (using measured declines in house prices at the onset of the Great Recession) and (ii) income losses (using data from the first months of the COVID-19 pandemic). Hereafter, we refer to these episodes as “aggregate” shocks, using quotation marks to remind the reader of the uneven distribution of these shocks over households with varying degrees of FD. Based on our findings, we assert that more FD means more consumer fragility, and—at least in the case of the collapse of house prices—greater macroeconomic

vulnerability.

More specifically, we develop a model that collects all the (zip-coded) geographies that fall in a given quintile of incidence of FD. The result is five artificial economies made up of the zip codes experiencing the relevant quintile of FD. Within each economy, our model of consumption is rich enough to encompass heterogeneity in income risk, life-cycle consumption needs, housing, debt repayment, and, importantly, non-repayment (delinquency) and formal default (bankruptcy). To ensure that we carefully incorporate heterogeneity present in each of the “regions”, we estimate the model separately for each of the five categories of zip-code-level FD. We then use this battery of estimated models to demonstrate the channels at work and show—via region-specific counterfactuals—that these channels are relevant for the localized, and then aggregate, consumption response to “aggregate” shocks.

Our analysis delivers two conclusions. First, the impact of “aggregate” shocks on regional spending is very unequal, where regions with higher prior FD are prone to suffer a more significant decline. But the difference is not just because “aggregate” shocks are unevenly distributed. It is also because these regions react differently even after receiving the *same* shock. Our counterfactuals show, for example, remarkably different consumption responses by region when we hit each region with a baseline 9 percent decline in house prices—to mirror data from Great Recession. In the areas with the lowest FD, consumption *increases* by almost 4 percent, as some households can buy houses at the now-lower prevailing prices, and then allocate the savings to more non-durable consumption. In contrast, in the regions with the highest incidence of FD, the same shock delivers the more familiar pattern of a sharp 1.1 percent consumption decline.¹

Our model suggests that, through these same mechanisms, the current macro shock to income arising from the COVID-19 crisis will also generate unequal responses of consumption across regions with different levels of FD. To evaluate this, we construct a counterfactual exercise where each region suffers an 8 percent decline in income (as opposed to the house price declines that were dominant in the Great Recession but do not characterize the more recent shock). We find that consumption declines across all “geographies”, including by 2.8 percent in

¹We corroborate this pattern with [Mian et al. \[2013\]](#)-style regressions in the appendix.

the region with the lowest FD, but declines far more—by 4.2 percent—in the region with the highest FD at the onset of the shock. In this clear sense, the presence of FD amplifies the consumption response to shocks that may hit any region.

What if we had ignored FD in our analysis? To address this, we consider two counterfactual economies that do not allow for unsecured debt repudiation, and hence do not feature anything like the persistent high-credit-cost state of FD. The results show that the economy’s response to the decline in house prices during the Great Recession would have been an *increase* in aggregate consumption, instead of the small overall decline we find in our benchmark economy with FD (which in turn varied across FD quintiles). Similarly, excluding FD from the model, the aggregate marginal propensity to consume *declines* by almost one-third—from 0.41 to 0.33—for the income shock meant to represent the COVID-19 shock. Our findings demonstrate, we believe for the first time, that FD may be a relevant aspect of the data to capture in analysis of the localized consumption response to “aggregate” shocks.

1.1 Financial Distress

Two formal definitions of FD are developed in [Athreya et al. \[2019\]](#) and are the same as used here: a significant fraction of available credit card credit is exhausted, as measured by the percentage of remaining credit, or there is delinquent debt.² In that work, both these measures of FD are shown to be relatively common (i.e., high incidences) and disproportionately accounted for by a smaller group of households persistently in FD. Thus, the empirics of FD in the U.S. suggest that individual consumption dynamics over longer-run periods are affected by distress for many, with some facing much more frequent difficulties.

On one level, FD resembles conventional measures of liquidity constraints. One definition defines FD this way: a household is in FD if it has exhausted more than 80 percent of its credit limit. Measures of indebtedness are also plausibly natural contributors to FD: given any fixed borrowing capacity, more debt means

²These two definitions are formalized for their use in this paper at the beginning of [Section 2](#).

less ability to handle the next shock that arrives. However, in reality, credit limits vary significantly across households, so our measure is better because it uses both debt and credit limits. Similarly, leverage could be an alternative to FD. Nevertheless, FD is broader primarily because it is defined to include information encoded in past debt repayment decisions, something done neither by current debt nor leverage. In that sense, FD may help identify households' characteristics such as attitudes toward debt and repayment, which are crucial to determine the consumption response to shocks.

Defined as we have done, FD offers an encompassing, easily measured, and timely way to gauge households' and the broader economy's vulnerability to shocks. It is encompassing because, unlike other measures, it does not require knowledge of the items on households' balance sheets or prices needed to compute measures such as net worth or leverage. For example, one may well have little measured wealth but substantial amounts of *poorly measured* wealth (e.g., cash in a mattress or, more often, assets with uncertain liquidation values) or access to supplementary credit from hard-to-view sources (e.g., family or business assets that can be liquidated). Similarly, individuals with low levels of observable net worth may not be constrained.³

By contrast, seeing an individual become significantly delinquent, or utilizing most if not all unsecured credit, is more telling. It is unlikely, given the costs associated with being delinquent or utilizing typically expensive unsecured credit, that there are hidden sources of cheap credit available or that the household seeks to increase its net worth position in preparation for retirement, and so on. More importantly, since the marginal cost of credit influences the marginal propensity to consume (MPC), and the latter is central to accounts of macroeconomic susceptibility to shocks, FD is a window into both the individual and the aggregate MPC. As for ease of measurement and timeliness, our measures of FD are built on rich and frequently updated credit bureau data at the individual level.

³Think of those in middle age who are beginning wealth accumulation for retirement. 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.

1.2 Literature review

Our work connects with several strands of the literature. At the most general level, our findings suggest that the inclusion of FD into macro models is important to capture the real options available to households seeking to avoid debt repayment and bankruptcy. Its inclusion allows our model to capture—when calibrated—rich available data that help the baseline model get closer into replicating realistic consumption responses to shocks. In this sense, our work is related to the contributions of [Kaplan et al. \[2014\]](#) and [Carroll et al. \[2017\]](#). While our focus is only on consumption, that interest is driven by the standard (old- and new-Keynesian) views that at high frequencies, what happens to consumption is vital for the determination of income. Thus, our findings also connect with emerging literature of heterogeneous-agent models with market incompleteness and new Keynesian features [[Kaplan et al., 2018](#)].

The most closely related papers in terms of the empirics we uncover are ones that also document the relationship between shocks during recessions and prior conditions. [Guvenen et al. \[2014\]](#) display the entire distribution of income losses across many recessions. Especially relevant to our work is their finding that income losses during a recession generally tend to be larger the lower a person’s pre-recession income. Our work also focuses on the link between prior conditions and recession outcomes in housing markets. Here, [Piazzesi and Schneider \[2016\]](#) similarly show that cheaper houses during the 2000s experienced a more significant boom-bust cycle than more expensive ones, using city-level data from Zillow. Finally, [Patterson \[2018\]](#) shows that regions with higher MPC faced more significant employment fluctuations during the Great Recession.

Our work complements and extends these papers. It complements the literature by explicitly showing the covariance between shocks during the Great Recession and the COVID-19 pandemic with households’ ex-ante FD. It extends the literature by assessing how this fact matters through the lens of a rich model that incorporates FD as a choice. In particular, it studies the necessary counterfactuals.

There is, of course, a larger set of related papers that emphasize the role of delinquency or bankruptcy for macroeconomic fluctuations. The main difference

between previous work and ours is that while we focus on FD before the shock arrives, those analyses study how allowing for skipping debt payments shapes the responses of macro variables. In particular, while [Herkenhoff and Ohanian \[2012\]](#) and [Herkenhoff \[2013\]](#) emphasize the importance of default for the dynamics of unemployment, [Auclert and Mitman \[2019\]](#) consider the Keynesian channels of aggregate demand (via sticky prices and aggregate demands externalities).

A more empirical group of papers use individual-level data to investigate the consumption response to a change in house prices. [Campbell and Cocco \[2007\]](#) focus on the differences between the life cycle and home ownership. [Aladangady \[2017\]](#) and [Aruoba et al. \[2018\]](#) obtain empirical results in line with our finding that greater FD is associated with higher MPC. Those papers use zip-code-level data to highlight the importance of household financial constraints in shaping consumption responses. We connect these findings to FD, emphasize the importance of the geographical distribution of FD and house price shocks, and use a life-cycle model to compute counterfactual exercises that help us understand the aggregate implications of shocks and FD.

The work of [Mian et al. \[2013\]](#), replicated and extended by [Kaplan et al. \[2016\]](#), is essential to acknowledge here because that paper is the central reference when it comes to examining the response of consumption to the decline in home prices across zip codes. Since [Mian et al. \[2013\]](#) also analyze how MPC varies with leverage, in our empirics we make clear that our results on the role of FD is *not* merely repackaging leverage.⁴

Our model is closer to the model of mortgage default developed by [Hatchondo et al. \[2015\]](#). However, it incorporates (i) default on secured and unsecured debt as in [Mitman \[2016\]](#), (ii) formal and informal default as in [Athreya et al. \[2017\]](#),⁵ and (iii) five regions, each represented by an heterogeneous-agent model as in [Dupor et al. \[2018\]](#).

Several other papers are related to our work because they use heterogeneous-agent models to analyze the decline in consumption after house price shocks or,

⁴We establish the difference between the two in Figure A3 of appendix Section A.3, and Table 1. Note also that in appendix Section 6, which includes a few regressions to test whether our model’s conclusions are sensible in the data, we control directly for housing leverage. Our results remain unchanged when we do so.

⁵See also [Athreya et al. \[2015\]](#) and [Athreya et al. \[2019\]](#).

more generally, during the Great Recession. [Berger et al. \[2018\]](#) was the first paper to study how prices affect consumption in a heterogeneous-agent model with incomplete markets. They show how consumption responses depend on factors such as the level and distribution of debt, the size and history of house price shocks, and the credit supply level. [Kaplan et al. \[2019\]](#) build a quantitative model with long-term mortgages and default. Their key new component is the change in expected house price growth, which help accounting for the joint evolution of house prices and consumption during the Great Recession. [Garriga and Hedlund \[2017\]](#) use a model of housing search to show that an endogenous decline in housing liquidity amplifies the decline in consumption during the Great Recession.

Finally, our paper is related to a rapidly emerging COVID-19 literature. Using a very different data source, [Chetty et al. \[2020\]](#) show that early during the pandemic, spending patterns declined sharply in sectors that require physical interaction, because of layoffs, particularly of low-income employees. Also related, [Kaplan et al. \[2020\]](#) document that individuals in vulnerable occupations have lower labor incomes and lower liquid wealth. Lastly, [Glover et al. \[2020\]](#) emphasize the different economic effects of the pandemic on young and old individuals.

Overall, this emerging literature is in line with our interpretation of the COVID-19 crisis as a shock to employment and earnings that disproportionately affect areas with higher FD. Of course, our analysis of the COVID-19 crisis is only complementary to the Great Recession analysis because this is an ongoing event, and the availability of information is rapidly growing.

The remainder of the paper is structured as follows. In Section 2, we lay out the key facts related to the geographic variation in FD in the U.S. and the way that this variation is correlated with housing wealth losses during the Great Recession and projected income losses during the COVID-19 pandemic. With those facts established, we turn in Section 3 to our model, which as stated above, is capable of incorporating the desired margins of adjustment—and the costs associated with making them. Section 4 presents the parameterization. Section 5 contains the results, and Section 6 offers concluding remarks.

2 Financial Distress and “Aggregate” Shocks: Three Facts

The empirics we develop below will make use of two main definitions of FD developed by [Athreya et al. \[2019\]](#). The first of these, labeled *DQ30*, is the percentage of individuals with a credit card account at least 30 days delinquent at some point during the year. The second measure, labeled *CL80*, is the percentage of individuals within a zip code who have reached at least 80 percent of their credit limit over the same time interval.⁶ Using these definitions, we demonstrate that (i) the incidence of FD varied substantially across geographies at the onset of the Great Recession, (ii) the size of the shocks that occurred during both the Great Recession and the COVID-19 pandemic varied substantially across geographies, and (iii) the incidence of FD prior to both the Great Recession and COVID-19 pandemic and the size of shocks experienced in an area were significantly positively correlated.

Fact 1: household financial distress is unevenly distributed across zip codes

FD, as we have defined it, provides a useful and timely indicator of the financial health of a zip code and is easily accessible (in our case, via Equifax data). Figure 1 shows that both of our measures of zip-code-level FD convey the same message: the incidence of FD varied widely in 2002, which we take to be early enough to describe FD conditions before the Great Recession.⁷ Indeed, no state can be characterized as having entirely high or low FD, though FD does seem to be highest in the Southeast and Deep South.

These national pictures mask a high degree of dispersion within individual cities. Take, for example, two contiguous zip codes in St. Louis, Missouri: 63110 to the north and 63105 to the south. In 2006, 6.8 percent of households in the the southern zip code were in FD (*DQ30*), while more than twice that portion,

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

⁷The wide variance in FD shown here is not unique to 2002; similar maps from other years up to the present day reveal the same.

(a) DQ30



(b) CL80



Figure 1: National Maps of FD Dispersion in 2002
Source: FRBNY Consumer Credit Panel/Equifax.

15.5 percent, were in FD in the northern zip code. When housing prices started to collapse in, the loss in home value as a percentage of net wealth varied substantially as well:⁸ the southern zip code lost 0.5 percent, while the northern zip code lost twice as much, 1.0 percent.

The starkly different experiences of these two adjacent zip codes in terms of FD and wealth loss is not an anomaly. In our sample, the standard deviation of FD using DQ30 across Metropolitan Statistical Areas (MSAs) is 0.024, but the average standard deviation of zip codes *within* a given MSA is nearly twice that, 0.045. Similarly, while the standard deviation of CL80 is 0.026 across MSAs, it is 0.053—again, roughly double that on average—across zip codes within MSAs. In sum, differences in financial distress within MSAs are larger than between MSAs. Intuitively, this spatial concentration appears entirely consistent with the more general spatial stratification by economic condition exhibited in most, if not all, U.S. cities.


Aside from purely geographic heterogeneity, the variation—and inequality—in FD can also be seen in the distribution of the *quantity* of debt in delinquency. Figure 2 presents the Lorenz curve for the distribution of (at least 30-day) delinquent debt. We see that the top quintile of debt holders hold more than half of this debt. This is true both for credit card debt and for total debt. Conversely, the bottom 40 percent of debt holders account for less than 20 percent of delinquent debt.

While the data displayed so far are cross-sectional, this snapshot of dispersion is indicative of long-term characteristics. Athreya et al. [2019] use data at the individual level to show that FD is remarkably persistent under similar measures. For example, conditional on being in FD today, an individual is roughly four times more likely to be in FD two years from now than the average person. FD is similarly persistent at a community level, and even more so than would be expected if individual-level FD persistence were the only factor at play.⁹

⁸The method that we use to assign home values and net wealth to zip codes is described in appendix Section A.2.

⁹This is documented in appendix Section A.3.1.

Figure 2: Distribution of Delinquent Debt across Zip codes, 2002




Source: FRBNY/Equifax CCP.

Fact 2: “Aggregate” shocks are unevenly distributed across zip codes

The Great Recession and the COVID-19 pandemic are both clearly macroeconomic, or “aggregate”, events. Nevertheless, they did not effect all households and geographies uniformly, but landed on each with greater or lesser severity. We now describe how we represent each event within our model.

Great Recession: In terms of the “initial” shock experienced by households during the Great Recession, it is clear that there was a massive decline in house prices that began before the recession and continued well after. This drop in prices substantially damaged household balance sheets on average; but in keeping with the theme of this paper, it did not do so with any sort of uniformity across the country. Indeed, in a non-negligible portion of zip codes, the median home value rose, as shown in Figure 3. More familiar, of course, is the action on the other tail of the distribution: in a non-negligible share of zip codes, the median home value fell more than 50 percent.

Figure 3: Distribution of Home Price Losses, 2006-2012




Source: Zillow.

COVID-19 Pandemic: Unlike a decline in wealth from a decline in house prices, as in the Great Recession, the current pandemic essentially placed a substantial “tax” on certain forms of consumption (e.g, restaurants), as well as on some types of production (e.g., meatpacking). This impulse then translated very quickly into a change in labor demand, and hence changes in employment and wages. Moreover, as people pursued social distancing to mitigate the spread of COVID-19, employment in specific industries was (far) more negatively impacted than in others. In particular, industries such as “Accommodation”, “Food Services and Drinking Places,” and “Arts, Entertainment and Recreation” experienced catastrophic losses, as their usual business models require public interaction considered dangerous during a pandemic. Conversely, other sectors were far less affected, with some even experiencing increases in activity, e.g., the home improvement and grocery sectors.

Figure 4 plots the data on employment for the leisure and hospitality subsectors. There were large declines during April, amounting to more than 40 percent for the subsectors combined. There has been some recovery afterword, but prolonged social distancing measures have prevented these sectors from attaining

a sort of “V-shaped” recovery. Significantly, during July, many states reinstated their lockdown measures, due to large increases in COVID-19 infections.

Figure 4: Change in Employment in the Leisure and Hospitality Sector




Note: Seasonally adjusted data from the BLS.

This uneven effect across sectors and workers translates into an uneven effect across the zip codes in which those workers live. Using Longitudinal Employer-Household Dynamics (LEHD) Origin-Destination Employment Statistics (LODES) data,¹⁰ we calculate the share of workers at the zip-code level whose primary job is in “leisure and hospitality.”¹¹ The resulting variation is very clear, as seen in Figure 5. A substantial share of zip codes feature employment that is much more concentrated than the national average, while a substantial portion of zip codes exhibit the reverse pattern. The absolute variation is also large by any meaningful metric: 1-2 percent of zip codes have at least 30 percent of their employment in this extremely hard hit sector.

¹⁰We use data from 2017, which is the most recent year of data available.

¹¹We identify workers in leisure and hospitality as those whose primary job is in either “Accommodation and Food Services” (NAICS sector 72) or “Arts, Entertainment, and Recreation” (NAICS sector 71). Importantly, we present the share of workers living in each zip code whose primary *job* is in one of the sectors, as opposed to the share of workers within one of these subsectors employed by businesses in a given zip code.

Figure 5: Distribution of Employment Shares across the Leisure and Hospitality sector



Source: Census LODES. The vertical dotted line correspond to household-weighted national means.


Fact 3: Ex-ante FD and the share of the “aggregate” shock hitting a region are positively correlated

Thus far we have established (i) that FD—at the zip-code level—varies widely in the U.S. and (ii) that plausibly exogenous shocks (house price shocks in the Great Recession and work cessation in the COVID-19 pandemic) have caused uneven consequences across areas. We now establish our third fact: FD is a relevant pre-existing economic condition. In other words, we show that the presence of FD *prior* to shocks that are clearly “aggregate” in their effect on national-level outcomes contains “news” about the severity of the shock when it does finally arrive. And the news is not good: in both the Great Recession and in the ongoing COVID-19 pandemic, the most financially distressed households (again, at the zip-code-incidence level) were hit the hardest.

Great Recession: Starting with the Great Recession, Figure 6 shows that home values during this event declined the most in more financially distressed

communities. By 2012, regardless of FD, median home prices declined on average by around 15 percent relative to their 2006 levels. However, home price declines in zip codes with higher FD were in many cases twice that, or worse.

Figure 6: Regional Changes in House Prices by Financial Distress




Note: FD is measured with DQ30, which is the share of individuals who are at least 30 days delinquent on a credit card at some point in a given year. For ease of viewing, the data have been divided into 40 bins with respect to DQ30, and each dot represents the mean of that bin weighted by the housing wealth in each zip code as of 2006.

Perhaps worst of all, households hardest hit were not diversified. Specifically, we find that households with high financial distress also tended to hold a larger share of their net wealth in their homes. This implies that when losses are measured as a *percentage of 2006 net wealth*, home value losses are even 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 with higher FD. Appendix Section [A.3.3](#) illustrates this relationship.

COVID-19 Pandemic Similarly, in the COVID-19 pandemic, the declines in hours worked and employment were systematically larger in more financially distressed communities. As shown in Figure [7](#), there is a strong and consistent positive relationship between FD incidence at the zip-code level (measured by

the incidence of DQ30 in 2018) and the share of those area’s workers employed in “Leisure and Hospitality.” We also include the “Retail and Trade” sector to show that another most-affected sector also displays a similar pattern.¹²

Figure 7: Share of COVID-19 “Affected” Employment by Financial Distress




Sources: Census LODS and FRBNY Consumer Credit Panel/Equifax. Each dot represents the mean of a DQ30 bin weighted by the number of households in each zip code.

A natural conjecture, then, is that income losses among high-FD areas will be more significant in percentage terms than those of low-FD areas. To investigate this, we complement this information with the Household Pulse Survey. There, the Census asked households if they experienced income losses. To estimate the relationship between the share of households not affected (with no income loss) and the incidence of FD, we leverage data from the 10th wave. That is, we use the Pulse Survey to calculate the state-level shares of individuals who report “no earnings losses since March 13, 2020 (for self or household member).” We merge these state-level responses with our preferred Equifax FD measure (DQ30). Figure 8 shows that states with a higher incidence of FD tend have a lower share of households that escape the COVID-19 shock altogether—i.e., who have no labor earnings losses since March 13.

¹²According to BLS data, retail trade employment fell by 15 percent in April and recovered to a year-over-year decline of 8 percent in June.

Figure 8: Share of COVID-19 “No earnings losses” by Financial Distress



Sources: Census Pulse Survey Wave 10 and FRBNY Consumer Credit Panel/Equifax. Each dot represents the average state-level response of “% reporting no earnings losses since March 13.” Dashed line represents line of best fit, weighting each state by population.

Overall, the facts presented in this section suggest areas with higher FD may be more severely affected by the ongoing COVID-19 pandemic and the efforts to contain it.¹³

We turn now to the development of a model aimed at delivering an understanding the role of FD in macroeconomic vulnerability. As will become clear, the model is a rich one: it takes household consumption seriously, including housing and the contractual arrangements—renting or buying—used to obtain it, and features secured and unsecured debt and debt default.

¹³In a [pair of graphs](#) frequently updated, we additionally document that communities with high FD seem also to have higher numbers of COVID-19 cases and deaths per capita. Falling ill can come with a host of repercussions including medical bills, lost time at work, and stringent quarantine instructions. To the extent that the severity of community responses is positively correlated with the local severity of the pandemic, this may also point to additional consequences including lengthened stay-at-home orders, strained public health resources, and stronger local preferences against engaging publicly with local businesses. We will not explicitly model any of these consequences. If we did, however, the result would be an income shock more strongly correlated with FD, magnifying the results that we do present.

3 A Life-Cycle Model of Housing and FD

The question we are interested in is straightforward: How does (an area's) household financial health, as measured by FD, matter for the transmission of housing and income shocks to consumption? Given that FD is partially endogenous, however, answering this question meaningfully requires a model of debt acquisition, debt repayment, and consumption decisions. We now lay out such a model. In subsequent sections, we deploy it to measure, via specific counterfactuals, the role of FD in the response of consumption to housing and income shocks, including a quantification of the importance of the positive correlation between initial FD and these shocks.

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_n , which depends on age n . Each agent works for a finite number of periods and then retires at age W . Critically, all agents are subject to risk in their income y (specified below). Lastly, agents will be allowed to differ in the rate at which they discount the future. Specifically, a share p_L of the population has a discount factor of β_L , while the remaining share has a discount factor of $\beta_H \geq \beta_L$.¹⁴

With respect to markets, households have (limited) access to credit and each period choose non-durable consumption c , housing h , mortgages m' , and financial assets (or debt) a' . Households may choose to obtain housing services through homeownership or by renting. These options are an important form of heterogeneity to incorporate ex-ante, given the differences observed in homeownership rates across income categories in U.S. data.

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,

¹⁴Heterogeneity in the discount factor is common in macroeconomics at least since [Krusell and Smith \[2003\]](#). However, the modeling and the calibration of β heterogeneity here follows closely [Athreya et al. \[2019\]](#).

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.¹⁵

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

With probability $1 - \eta$, then, a household's rolled-over debt is not discharged. In this case, the household pays a "penalty" rate, r_R , of interest higher than the average rate paid by borrowers.¹⁶ Moreover, in any period of delinquency, we prohibit saving, and since the agent did not borrow but failed to repay as promised, their consumption equals income. Second, as in standard models of unsecured debt, agents may invoke formal default via a procedure that represents consumer bankruptcy (BK). If this is the path chosen, all debts are erased, and in the period of filing for bankruptcy, consumption equals income net of the monetary cost f of filing for bankruptcy.

3.2 Nonhomeowners

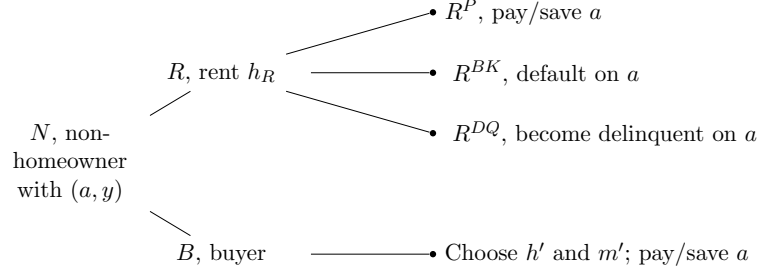
The options faced by a nonhomeowner with assets a and income y are represented in Figure 9. First, they can choose to either rent or to buy a house and become a homebuyer. If renting is chosen, the nonhomeowner must decide between the three options described below. There is a letter associated with each 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

¹⁵Housing choices, mortgages, and foreclosures are modeled as in [Hatchondo et al. \[2015\]](#).

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

variable a and y is N . For the sake of brevity, our formal description of this recursive problems is presented in Appendix B.

Figure 9: 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

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

Here, γ 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

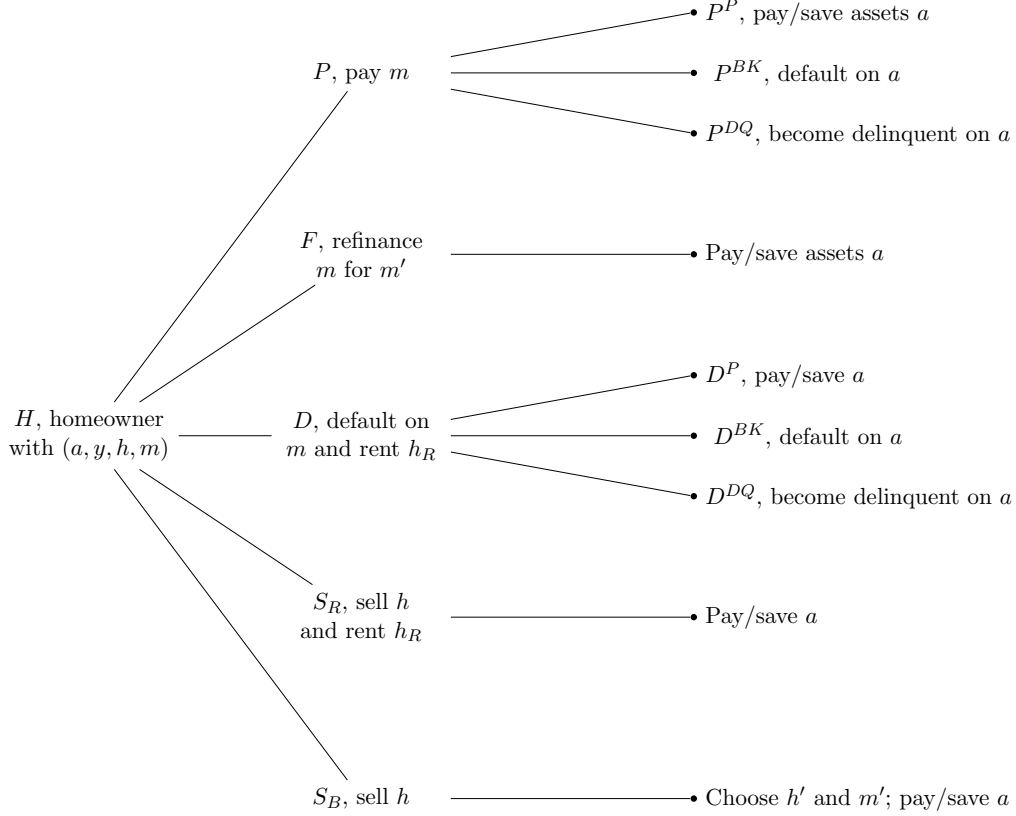
The choices available to an existing homeowner are presented in Figure 10. A homeowner's problem is more complex. On the financial asset dimension, homeowners must decide to default or repay their unsecured debt. On the housing dimension, homeowners can (i) pay their current mortgage, (ii) refinance 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 10: 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

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

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 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 renter who is also delinquent on existing debt. This means that consumption is given by $c = y$ and financial assets tomorrow are

equal to

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

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 B.

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

given by:

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

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

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

while, 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 $(1 - \gamma)$. 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 Estimation of the model to capture five FD “regions”

In order to most closely tie our empirical and quantitative work together, we need to take a stance on what a geographical region means in the model and data. A crucial feature is that even inside a zip code, we would need a force to deliver heterogeneous outcomes across agents to capture the fact that in *any* zip code, only a fraction of households are in FD. Defining a region as a zip-code, county, or even state would be computational prohibitive, as it would require a large number of estimations of our baseline model.

Thus, as a balance between expanding the reach of the model into more granular data and preserving practicality, we proceed as follows. First, we order the zip-codes in our sample by their incidence of FD and split the data into quintiles (5 groups, each with the same population size). Next, we construct five “regions” that *combine all zip codes that fall within each given quintile of FD*. We then treat these as “economies” or “geographies”, calculate several statistics (e.g., FD, income, wealth, and homeownership rate) for each region, and use these moments as targets for five different estimations of our baseline model. The statistics obtained are shown in Table 1.

By construction, FD is increasing across quintiles, and in terms of the absolute levels of FD (as defined by DQ30), we see that it increases from 8.6 percent of households in quintile 1 (Q1) to nearly triple that (23.5 percent) in quintile 5 (Q5). This is a first, and clear, indication that people in different quintiles tend to be differently positioned when it comes to their balance sheets.

Table 1: Descriptive Statistics by Quintile of DQ30 in 2002

	Quintiles of DQ30 in 2002				
	1	2	3	4	5
	<u>Wealth</u>				
Income Per Household (HH) \$000	91.75	65.26	53.51	46.22	39.86
Net Wealth Per HH. \$000, ages 25-55	358.5	216.0	164.5	127.1	88.12
Fin. Wealth Per HH \$000, ages 25-55	321.9	201.4	154.6	123.4	83.00
Net Fin. Wealth Per HH \$000, ages 25-55	224.0	128.1	95.00	72.71	42.13
Median Home Value \$000	297.0	219.0	179.9	154.8	128.6
	<u>Human Capital</u>				
Less Than High School	7.659	11.95	16.69	19.63	23.73
High School	19.70	24.78	26.82	27.99	29.23
College	72.64	63.27	56.49	52.37	47.04
Age	44.27	43.61	43.27	42.84	42.64
	<u>Debt</u>				
Percent of HHs that Own a Home	76.30	71.93	68.76	64.25	61.69
Percent of HHs with Housing Debt	49.77	44.67	39.83	36.27	31.84
Housing Debt per Home Owner \$000	135.0	102.3	83.91	73.38	58.95
CC Debt Per Household \$000	5.238	4.803	4.407	4.171	3.806
Housing Leverage	44.11	47.98	44.57	46.04	43.36
	<u>Delinquency</u>				
HHs with housing debt and in FD / HHs (in %)	5.910	8.555	10.82	13.32	19.46
HHs with housing debt / HHs in FD (in %)	33.31	30.72	28.37	26.90	25.99
Foreclosure Rate	1.520	1.812	2.239	2.579	3.335
Bankruptcy Rate	0.392	0.553	0.631	0.648	0.639
DQ30	8.566	12.11	14.92	17.83	23.54

Note: Here, housing debt refers to a mortgage or home equity line of credit. Housing leverage is measured as housing debt divided by the total housing wealth in each geography. The number of households weights all means, except housing debt per homeowner, which is naturally weighted by homeowners. “ages 25-55” signifies that for the corresponding rows, we used financial wealth aggregates from the SCF for individual from 25 to 55 years old. This is done because elderly populations hold a large share of financial wealth, and our model economy is calibrated for individuals 25 to 55 years old.

Naturally, FD is inversely related to various other measures of economic health, wealth, and human capital. Areas with high FD tended in 2002 to have lower incomes, net wealth, and home values. Lower wealth in high FD areas prevents these areas from sustaining higher levels of debt, both in terms of housing debt and, perhaps more surprisingly, credit card debt. This lower credit card

debt arises because despite zip codes with high FD using a higher proportion of their available credit, they also tend on average to have significantly lower credit limits. On the other side, zip codes with low FD enjoy the double bonus of having a high credit limit and having used a lower portion of that limit. Thus, from an ex-ante perspective, the latter is better situated to weather financial losses. In terms of human capital, people in the highest FD quintile are less than half as likely to have earned a high school diploma as those in the lowest FD quintile.

Since we intend to look at the interaction between FD and housing shocks, and since those in high-FD zip codes are somewhat less likely to own homes, it would be problematic if the differences in FD across zip codes are driven mainly by people who do not own homes. To examine this, we need to identify at the individual level homeownership and FD, something we cannot do with Equifax. We proxy for homeownership within the Equifax data by using natural objects that we can observe: whether an individual has either a mortgage or a home equity line of credit (housing debt).¹⁷ The bottom panel of the table shows that when we consider the fraction of people identified to both own a home and be in FD, the resulting differences between quintiles are similar in magnitude to those of FD considered directly. Taken as a whole, this is important, as it clearly suggests that it is highly unlikely that the dispersion in FD is being driven by people who do not own homes.

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 the statistics mentioned above for each of the five regions.

4.1 Assigning first-stage parameters

Table 2 collects the parameters set externally. A period in the model refers to a year. Households enter the model at age 25, retire at age 65, and die no later than age 82. We set the risk-free interest rate at 3 percent. In addition, we

¹⁷Of course, this method does not allow us to identify homeowners who have completely paid off their homes and have no home equity lines of credit. The percent with housing debt usually underestimates the percentage of households that own the home they live in by about a third.

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.

Turning to preferences, we make two data-disciplined changes to an otherwise standard formulation. First, as previously mentioned, we follow [Athreya et al. \[2019\]](#) and assume 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 $\beta_L \leq \beta_H$. Let s_L denote the share of the population of type L. This allows the model to capture well the joint distribution of net financial wealth, delinquency (incidence and persistence), and bankruptcy. Second, it matters that our model match as well as possible the joint distribution of homeownership and FD. Here, we find that a simple allowance for the “specialness” of owner-occupied housing (presumably capturing a variety of benefits that ownership confers) relative to renting helps reconcile theory and data. This is represented in a simple manner: the utility u derived from consumption c and from living in a house of size h displays a constant elasticity of substitution between the two goods:

$$u(c, h) = \frac{((1 - \theta)c^{1-1/\alpha} + \theta(1 + \theta_i^R \mathbf{I}_{\text{renting}})h^{1-1/\alpha})^{(1-\gamma)/(1-1/\alpha)}}{1 - \gamma},$$

where γ denotes the risk aversion parameter, α governs the degree of intra-temporal substitutability between housing and non-durable consumption goods, and θ determines the expenditure share for housing. The parameter $\theta_i^R \mathbf{I}_{\text{renting}}$ captures the type-specific ($i \in \{L, H\}$) disutility from renting relative to owning a house. Following [Hatchondo et al. \[2015\]](#), we set γ to 2, α to 0.5, and θ to 0.11. Since we ultimately calibrate the rental house size h^R to match each region’s homeownership rate, we normalize the value of the disutility of renting for individuals with a low discount factor, $\theta_L^R = 0$. Thus, what remains to be determined is the region-specific value of θ_H^R .

Following [Livshits et al. \[2007\]](#), the penalty rate for delinquent debt is set at 20 percent annually and the bankruptcy filing costs are at 2.8 percent 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.

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 percent to 14 percent, with an average replacement rate of 47 percent. The age-specific survival probabilities follow [Kaplan and Violante \[2010\]](#).

4.2 Estimating the remaining parameters

The remaining parameters to be determined are (i) the discount factors of impatient types β_L , (ii) the discount factors of patient types β_H , (iii) the share of impatient types in the population s_L , (iv) the probability of delinquent debt being fully discharged η , (v) the house price per unit p , (vi) the rental house size h^R , and (vii) the disutility that type-H agents receive from renting versus owning a house θ_H^R . We estimate these seven parameters so that model-simulated data replicates some critical features of the data about homeownership, financial wealth, and FD for each of the five regions we construct.

Table 3 presents the model's fit for each of the quintile-specific moments.

Table 2: 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	Survey of Consumer Finances 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/ average income	Livshits et al. [2007]
r_R	0.2	Roll-over rate on delinquent debt	Livshits et al. [2007]

The model does an excellent of matching differences in financial wealth across the five regions. Additionally, it replicates the fact that homeownership declines as regional FD rises and does a good job of matching the share of individuals in FD that have housing debt. 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. One shortcoming of the model is that it struggles to precisely match the ratio of median home values to mean income. While in the data there is no systematic pattern of this ratio with FD, the model suggests the ratio declines slightly as FD rises.

The rest of the table focuses on FD and shows that the model does a good job of matching the overall regional patterns as well. Indeed, the model nearly exactly matches the fact that average delinquency rates rise with each quintile of FD, and so do bankruptcy rates. Additionally, the model matches the fact that the persistence of FD actually falls as the quintile number increases.

Table 4 shows the resulting parameter estimates and reveals some systematic differences across the quintiles of FD. Most notably, the share of impatient people systematically rises across the quintiles. For example, in Q1, less than a quarter

Table 3: Regional Calibrations

Moment	Q1		Q2		Q3		Q4		Q5	
	Data	Model	Data	Model	Data	Model	Data	Model	Data	Model
Wealth / Income	2.49	2.46	2.01	1.97	1.79	1.78	1.62	1.56	1.09	1.09
Homeownership rate	76.9	80.6	72.7	75.0	68.9	70.3	65.3	68.1	61.6	61.0
Home value / Income	3.23	3.37	3.30	3.11	3.42	3.13	3.33	3.07	3.23	3.08
DQ rate (in %)	8.6	8.3	12.1	11.0	14.9	13.8	18.9	17.3	23.6	23.3
BK rate (in %)	0.37	0.38	0.54	0.54	0.62	0.58	0.64	0.67	0.64	0.64
Persistence of FD	5.45	5.23	4.74	4.41	4.06	3.84	3.55	3.61	2.88	2.76
With housing debt / In FD (in %)	33.6	31.3	31.3	30.9	28.6	27.0	27.2	26.7	26.0	25.5

Note: “Wealth/Income” represents mean net financial wealth divided by mean income; “Home value/Income” is the median home value divided by mean income, and “With housing debt / In FD” is the percent of the population with housing debt, conditional on being in FD.

of the population discounts the future relatively more. In contrast, in the fifth quintile, over half of the population is impatient. In terms of the values for the discount factors, the model requires only modest differences across quintiles but large differences across types. For example, the high discount factor is essentially identical across the quintiles of FD, and the low discount factor β_L is only significantly lower in the Q5 of FD compared to the other four.

Lastly, the data—filtered through our framework—imply significant differences in the utility of rental house sizes between types, θ_H^R , regardless of the quintile of FD. One way to interpret the parameter θ_H^R is that rental houses are perceived to be of different sizes by agents of different types. For example, in Q3, the coefficient of 4.38 implies that β_H -type households perceive rental houses as about 20 percent of the size perceived by β_L -type households. This difference allows the model to match the low homeownership rate among households in FD (mostly β_L types)—approximated by the percent of households with housing debt among those in FD—together with a high overall ownership rate.¹⁸

¹⁸Indeed, with a single parameter governing the size of rental houses, the model-implied ownership rate is biased away from the data value and has the wrong FD composition. The comparatively high ownership rate of low-FD individuals dictates a small rental house size, but with a single parameter, even high-FD individuals own houses and so the overall ownership rate is too high. Conversely, the comparatively low ownership rate of high-FD individuals dictates a large rental house size, but with a single parameter even low-FD individuals choose to rent, making their overall ownership rate counterfactually low. Allowing for differences in rental house size by β type resolves this tension.

Table 4: Regional Parameter Estimates

Parameter	Q1	Q2	Q3	Q4	Q5
Low discount factor β_L	0.64	0.65	0.62	0.56	0.51
High discount factor β_H	1.04	1.03	1.03	1.03	1.02
Share pop. w/ low discount factor s_L	0.24	0.30	0.35	0.39	0.50
Rental house size h^R	4.02	3.81	3.58	2.92	2.50
Utility of renting versus owning for H -type θ_H^R	7.28	5.63	4.38	4.12	14.2
Owner-occupied house price p	3.63	2.74	2.36	2.45	2.51
Discharge prob. of DQ debt γ	0.87	0.75	0.70	0.68	0.67
LTV λ	0.9	0.9	0.9	0.9	0.9
Average earnings (relative to Q3)	1.41	1.17	1.0	0.89	0.76

5 Quantitative Exercises

We now use the model to understand the relationship between financial distress, shocks (to housing wealth, then income), and the response of consumption during the two macroeconomic events we consider. This analysis requires, first of all, that we generate within the model a stylized Great Recession and then an episode that captures some key aspects of the COVID-19-induced lockdown. In our quantitative analysis, both shocks will be exogenous. Of course, house prices and labor income have endogenous components (see, e.g., [Garriga and Hedlund \[2017\]](#) for a rich analysis of the former, and of course countless business analyses of the latter). Our goal is not provide an account of these price movements, but rather to understand how *shocks* are unequally distributed and unequally transmitted into consumption.

We then use the model to uncover the micro-level mechanisms at work in an aggregate shock.¹⁹ We stress that our work is not an attempt to analyze the economic impact of the COVID-19 pandemic per se. Indeed, any macroeconomic shock for which we had relatively granular measures of idiosyncratic incidence would do; the two shocks we utilize are both recent and sizeable and between them cover two kinds of economic stress (net worth and labor income) that are empirically relevant. An additional impetus for using the ongoing pandemic in particular is that as a major macroeconomic event, it is relatively clean in its

¹⁹A validation of the primary mechanism is presented in the last part of the paper.

(extremely) exogenous nature, at least at the outset.

As just noted, the Great Recession and the COVID-19 pandemic attacked different parts of a household’s financial well-being. In the case of the Great Recession, household net worth was destroyed, while in the pandemic, income-generation effectively became impossible for a subset of households, certainly in the short-run. Thus, by studying both, we expand the reach of our analysis of how FD matters for macroeconomic outcomes.

5.1 A first aggregate shock: A collapse in asset valuations

A central aspect of the Great Recession was a large drop in home prices. We therefore replicate this event in our model by subjecting each of our calibrated “regions” to exogenous changes in house prices. Again, we remind the reader that our approach is to treat those in a category (specifically, quintile) of FD, gauge their response to a shock, and then compare this response to those of the other quintiles of FD. One aspect of our representation of the shocks is that they respect the data we presented in Section 2. Namely, that the shocks landed most heavily on areas that exhibited greater financial distress at the outset.

A key finding from these experiments is that our model implies very different consumption responses across “regions.” We find also that much of these differences remain *even when we subject the regions to the same shock*. Differences in initial FD alone appear to drive very disparate regional outcomes for a given shock. That is, FD matters.

Turning to details, we proceed in this part of the analysis by subjecting the stationary distribution of each region to an exogenous and unanticipated (but permanent) house price decline. Importantly, we allow for region-specific house prices shocks that mimic the previously documented house-price declines across different FD regions. To use the data presented in Section 2, we summarize the information into the five “regions” created. Because the model is yearly, we need a yearly change in house prices for each region. We selected the change between 2007 and 2008.²⁰

²⁰It is useful to note that we obtain very similar results using the average yearly change

The first row of Table 5 shows the shocks hitting the economy. The baseline decline in house prices is significantly uneven across FD “regions”: it is only 7 percent for Q1, but reaches 11.5 percent for Q5. The implied aggregate implications are presented in the sixth column. Note that the aggregate decline in house prices is 9.1 percent. The last column shows a counterfactual aggregate economy in which each region has a decline in house prices of 9.1 percent.

The rest of Table 5 shows the implications of the decline in house prices. Because the house price shocks are modeled as permanent changes, all the values presented are measured as percentage change relative to old steady-state averaged over three periods as in Dupor et al. [2019].²¹ The aggregate decline in consumption is only 0.03 percent. In terms of a MPC out of a change in house prices, this change implies that consumption declines less than 1 cent per dollar decline in house prices. To put this in context, Mian et al. [2013] estimate an MPC for nondurable spending of 1.6 cents per dollar and essentially zero for grocery spending.

Given our aim to understand the manner in which FD affects the ability of households (and by extension the macroeconomy) to withstand shocks, it is essential to focus on the change in consumption across regions. The contrast across quintiles is very striking: consumption *increases* by 0.42 percent in Q1 but decreases 1.32 percent in Q5. The changes in other variables offer clues about the mechanism. Note, for example, that household financial assets decline across quintiles after the shocks. Perhaps the most important is the change in unsecured debt, which declines by 14.5 percent for Q5 but only by 2.1 percent for Q1.

These differences across “regions” are relevant because they show that the response to “aggregate” shocks may be very different. However, they are also meaningful because they have aggregate implications. Comparing the aggregate results presented in the last two columns, we can see that aggregate consumption declines only slightly (0.03 percent) with the actual distribution of shocks and increases significantly (0.88 percent) with the shocks distributed uniformly.

between 2006 and 2009 as well.

²¹For example, if the change measured relative to the steady state is 2 percent and is preceded by a path of 2 percent in the first period, 3 percent in the second, and 4 percent in the third one, the change presented in the table would be 3 percent.

Table 5: House Price Shock Experiments


% chg in	Unequal shocks					Equal shocks	
	FD “Regions” or Quintiles					Aggregate	Aggregate
	Q1	Q2	Q3	Q4	Q5		
House prices	-6.99	-8.60	-10.0	-10.9	-11.5	-9.10	-9.10
Consumption	0.42	0.04	0.10	-0.37	-1.32	-0.03	0.88
Fin. assets	-1.63	-1.39	-1.73	-1.53	-1.50	-1.42	-1.41
Unsec. debt	-2.11	8.27	-4.47	-11.6	-14.5	-8.50	26.3
Home equity	-6.71	-9.39	-11.0	-12.6	-14.5	-10.1	-10.1
Ownership	0.38	-1.62	-1.55	0.44	0.09	-0.53	-0.53

Note: All values are measured as percentage change relative to the old steady state, averaged over three periods during the transition to new steady state.

In what we have reported so far, we have used the data directly, inclusive of the covariance structure summarized in “Fact 3” above. However, it is important to provide some isolation of how the distribution of the shocks across FD “regions”, purely on its own, works to alter the transmission of a shock. We therefore examine next the case with shocks identical across all “regions.” We see that total spending (consumption), i.e., spending aggregated across FD quintiles, *increases* by almost 1 percentage point. To understand where that difference in the aggregate numbers is coming from, Figure 11 presents the implied regional consumption responses from both experiments. In the case with equal shocks among “regions,” the differences between the most and least distressed regions are even more stark. In this case, we see that the least financially distressed region sees consumption increase by nearly 4.0 percentage points. This is the basis for our claim that relatively microeconomic, i.e., zip-code level, FD matters for who bears the burden of macroeconomic risk, and to some extent for macroeconomic vulnerability itself.

While the previous analysis helps illuminate the importance of accounting for regional heterogeneity in FD and the role of uneven shocks, it does not fully delineate the importance of modeling FD. To address this, we now conduct two more exercises. First, we consider what happens in a setting with no possibility of informal default or bankruptcy (referred to as “no FD”). This case is, in one sense, the standard case studied in most models of consumption, where neither formal nor informal default are typically allowed. This baseline case of course

Figure 11: Consumption Responses to House Price Shocks by Quintile of FD and Experiment Type



Notes: All consumption changes are relative to old steady states and are averages over three periods during the transition to new steady state. Dashed line represents economy-wide average of corresponding variable under benchmark case. Dotted line represents economy-wide average of corresponding variable under same shock case.


cannot make any contact with empirical notions of FD and also (really, hence) implies that all financial debt is risk-less because it is always repaid. Following that, we address the importance of modeling FD in a second counterfactual where we disallow *unsecured* borrowing altogether—think of this case as adding a zero borrowing constraint. We refer to this as the “no borrowing” case. Figure 12 presents the results of these two counterfactual scenarios.

Across the economies, the availability (or lack thereof) of FD matters substantially for the response of consumption to house price shocks. In general, across all economies, removing the possibility of FD shrinks the drop in consumption (or increases the jump). Focusing on the most distressed quintile/region/economy (Q5), the economy without FD has a minimal consumption drop instead of the -1.3 percent decline in the benchmark economy. In the economy with no unsecured borrowing, the difference relative to the benchmark is even more striking, with consumption in Q5 increasing by half a percentage point due to the decline in house prices.

The horizontal lines in Figure 12 represent the aggregate changes. The aggregate numbers reflect what happens in each “region”: removing FD reduces the drop in consumption (or increases the jump). In both counterfactual cases, the aggregate change in consumption would be positive instead of slightly negative. Again, this highlights the role of FD for the response of the economy to “aggregate” shocks.

The key mechanism behind the differential consumption response of the non-FD economy to the baseline reflects a well-known feature of models of defaultable debt: unsecured borrowing is risk-less, and consumers can borrow much more to smooth consumption. In the no-borrowing economy, precisely to deal with the inability to borrow, agents generally have higher asset positions to smooth housing shocks. Additionally, the richness of our model allows us to capture a more subtle effect running from wealth to consumption: better asset positions reinforce the income and substitution effects of lower house prices for “soon to be owners”, who can raise their non-housing consumption.

Figure 12: Consumption Responses by Quintile of Financial Distress and Debt Arrangement



Notes: All consumption changes are relative to old steady states and are averages over three periods during the transition to new steady state. The horizontal lines represent the economy-wide average consumption drop in each case.

5.2 A second aggregate shock: Income Loss

The economic downturn generated by the COVID-19 pandemic has (aside from its direct health effects) affected households differently than did the asset-price collapse that was the first manifestation of the Great Recession. Section 2 showed that areas with greater FD also tend to have larger employment shares in industries that were more affected by the social distancing that accompanied the COVID-19 pandemic. It also showed that areas with more FD had fewer households that were not affected by this crisis. In this section, we map these numbers into the five quintiles, or “regions,” for use in our calibrated model. We then present a simulation exercise in which agents are subject to earnings shocks that are obtained from the information in Section 2 and from professional forecasts for income losses over the rest of the year. Much like the results from the previous section, we find that the dispersion in income shocks results in significant dispersion in consumption responses across the FD distribution. Additionally, and

again in line with the previous section, we find that even if income shocks are uniformly distributed across the FD distribution, there remain notable differences in consumption responses across the quintiles of FD.

There are three steps in our calibration of the COVID-19 shock. First, we compute the share of workers in “leisure and hospitality” for each quintile. These are the *most affected* workers, as they have a decline in yearly earnings of 30 percent. The decline in income for this group was calibrated to be in line with the data presented in Figure 4.²² The findings are presented in the first column in Table 6. It shows the share of households in the “leisure and hospitality” sector increases across FD quintiles, ranging from 9.17 percent in Q1 to 12.1 percent in Q5. This fact naturally replicates the correlation of FD and the share of workers in this sector at the zip-code level shown in Figure 7.

The second step in our calibration of the COVID-19 shock is to identify workers who are not affected by the lockdown, perhaps because they can work remotely. Figure 8 already showed that the share of households that are not affected is decreasing in the incidence of FD at the state level, according to data from the Census Pulse Survey. We use that relationship to obtain the predicted share that is not affected for each of our “regions” and present the results in the second column of Table 6. While 52.2% of households in Q1 did not have income losses over the first 4 months of the lock-down, this share is only 38.4% for Q5.

Table 6: Distribution of COVID-19 related earnings losses by FD.

FD Quintile	Share of workers by Sectors:			Average Δ Income
	Most affected	Not affected	somewhat affected	
1	9.17	52.2	38.7	-6.26
2	10.0	48.9	41.0	-6.74
3	10.6	46.4	43.0	-7.09
4	11.1	43.7	45.2	-7.43
5	12.1	38.4	49.5	-8.13
Δ Sectoral Income, %	-30.0	0.00	-9.08	-
Implied Aggregate % Δ	-3.12	0.00	-3.88	-7.00

The third and last step is to derive how affected households are that belong

²²This provides a conservative estimate for the rest of the year given that the loosening of the lock-down during June generated a new wave of cases.

neither to the most-affected nor to the unaffected sectors. We set this number such that the aggregate decline in income for our 2020 simulation matches the predicted decline of 7 percent by the consulting firm Macroeconomic Advisers. Given that the predicted aggregate decline is so severe, and that about 40 percent of households will be unaffected, we need a decline in income for the “somewhat affected” households of 9.1 percent to finalize our calibration.

The last column in Table 6 summarizes the shock presenting the decline of average income by quintile. In the “region” with the lowest incidence of FD (Q1), the decline in income averages 6.26 percent, while the decline in average income for the region with the highest incidence of FD (Q5) is 8.13 percent. Next, we hit each of our five calibrated “regions” with income shocks following the distributions displayed in Table 6. These shocks are modeled as unanticipated and transitory. For instance, in Q1, we randomly select 9.17 and 38.7 percent of the households and reduce their incomes by 30 percent and 9.08 percent, respectively.

Table 7 shows the results. The first row repeats the information about the change in income for each “region.” Across quintiles it is clear that there are differences in how much consumption responds to income shocks. While individuals in Q1 decrease their consumption on average by just under 2 percent, individuals in Q5 decrease their consumption by more than 4 percent. These striking differences show that without any policy interventions, households were expected to respond very differently to the COVID-19 shock. As the sixth column shows, aggregate consumption would decline by 2.79 percent. Given the size of the shock, this implies a MPC equal to 0.33, which is comparable to MPCs out of transitory income shocks reported in Table 1 of [Carroll et al. \[2017\]](#) when looking at horizons of one year. Additionally, it is worth highlighting that since we distinguish between financial wealth and housing wealth when calibrating our economies, our model-based MPCs will in general be higher, as shown in [Carroll et al. \[2017\]](#).

Table 7 also presents information about other statistics that are useful to understand the differences across quintiles. One fact to highlight is the increase in bankruptcies, which is much higher for Q5 than Q1.

The last column in Table 7 shows how the aggregate change would look if the share of households in each group were the same for all quintiles. The results

Table 7: COVID Income Shock Experiments

% chg in	Unequal shocks					Aggregate	Equal shocks
	FD “Regions” or Quintiles						Aggregate
	Q1	Q2	Q3	Q4	Q5		
Income	-6.26	-6.74	-7.09	-7.43	-8.13	-7.00	-7.00
Consumption	-1.99	-2.35	-2.90	-3.20	-4.20	-2.79	-2.75
Fin. Assets	-1.45	-1.81	-1.88	-1.97	-2.33	-1.76	-1.78
Fin. debt	7.56	9.34	9.59	7.20	5.54	7.42	7.17
Home equity	-0.48	-0.67	-0.69	-0.77	-1.08	-0.67	-0.67
DQ incidence	11.87	11.48	9.18	7.37	7.33	8.88	8.70
BK incidence	4.07	10.35	12.72	13.75	14.33	11.68	11.01
Ownership	-0.69	-0.97	-1.23	-1.56	-2.08	-1.26	-1.23

Notes: Income shocks are transitory. All values are measured as percentage change relative to old steady state. Change in income and consumption are measured in period of shock.


Changes in all other variables are measured in period after shock.

indicate that the effect would be less severe, but changes in all the variables are small. For instance, the decline in consumption is now 2.75 percent versus a decline of 2.79 percent in the benchmark. Similarly, debt increases by more in the benchmark, but the difference is 7.42 percent vs. 7.17 percent.

However, the preceding finding does not mean that differences in the incidence of FD does not matter. Just as we did in the case of the shock representing the Great Recession, Figure 13 answers the question of what the consumption responses would look like if all quintiles were subject to the same distribution of income shocks. As can be seen, most of the behavior remains the same as in the benchmark. In both cases, there are significant differences in consumption across regions/FD quintiles. With the same shocks, the decline in spending by the least financially distressed is 2.2 percent, while it is about one-and-half times as large as that amongst the most financially distressed (3.6 percent). Another way to think about this finding is the following: If the entire United State were to look like our regions with the highest FD, the aggregate decline in consumption would be 3.6 percent instead of 2.8 percent, or roughly \$200 billion greater on impact.

Lastly, to uncover the role of FD as a mechanism in dealing with this aggregate shock to income, as opposed to asset valuation, we proceed as before and conduct the same pair of counterfactual exercises. To remind the reader,

Figure 13: Consumption Responses to COVID Income Shocks by Quintile of FD and Experiment Type



Notes: All changes are measured in period of shock. Dashed line represents economy-wide average of corresponding variable under benchmark case. Dotted line represents economy-wide average of corresponding variable under same shock case.


these counterfactual scenarios are (i) an economy where borrowing is allowed but default (and hence FD) is not, and (ii) an economy where debt is disallowed altogether. Figure 14 presents the results.

We see from the figure that both removing the FD option and removing the option to borrow are associated with significantly smaller aggregate consumption declines. In particular, note that aggregate consumption declines 2.3 percent instead of 2.79 percent.

6 Empirical validation

Our model suggests that FD matters: microeconomic distress is related to greater sensitivity to macroeconomic shocks. This was seen throughout the results of sections 5.1 and 5.2 where at an aggregate level, higher FD was associated with larger consumption declines in response to shocks. In our model, at the individual level, agents cut their consumption more drastically not just because FD prevents them from having access to credit, but also because other characteris-

Figure 14: Consumption Responses by Quintile of FD and Debt Arrangement



Notes: All changes are measured in period of shock. The dashed line represents the economy-wide average consumption drop under the benchmark model. The dotted line represents the economy-wide average consumption drop when no FD is allowed. The dash-dotted line represents the economy-wide average consumption drop when no borrowing is allowed.

tics correlated with FD. Using the case of homeowners as an example, those in FD mostly turn out to be those with higher discount factors (the “impatient” types) who in turn often have long histories of facing high borrowing costs in the unsecured credit market. As a result, their consumption is mainly financed through other means such as mortgage refinancing. When housing and income shocks arrive, these means vanish and they respond by aggressively cutting consumption. To what extent can additional evidence be brought to bear to validate this mechanism?

While we lack sufficiently detailed data at the individual level to corroborate this mechanism directly, in the case of the house price shock, we can test this result by asking what happens at a more aggregate level. That is, we can ask whether consumption in regions with higher FD actually responds more to housing price shocks. We argue that the answer is “yes.”

To this end, we now estimate the MPC out of housing shocks following the seminal work of [Mian et al. \[2013\]](#). In particular, we want to determine whether

MPCs vary in a significant fashion by FD holding constant other regional features such as income, wealth, etc. Formally, we estimate regressions of the form:

$$\Delta C_t^i = \alpha + \beta_1 \Delta HV_t^i + \beta_2 FD_t^i + \beta_3 (\Delta HV_t^i \times FD_t^i) + \beta_4 X_t^i + \epsilon_t^i. \quad (1)$$

Here, ΔC_t^i is the dollar change in consumption in geographic region i between t and $t+1$; ΔHV_t^i is the change in home value; FD_t^i is the level of FD in region i at time t ; X_t^i is a vector of other regional covariates that can be both in levels and changes; and ϵ_t^i is the error term.²³ The coefficient of central interest is β_3 , the *interaction* between FD and housing shocks. We focus on new auto purchases, as our measure of consumption at the county level. In terms of timing, all initial levels are measured in 2006 (except for FD which is measured in 2002), while all changes are measured between 2006 and 2009.

Table 8 reports the results of estimating equation (1). 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.

It is easiest to interpret the interaction term coefficients with some examples. Figure 15 shows how the coefficients in Column (2) of Table 8 translate into differing MPCs by level of FD. The dark set of bars represent the average MPC out of a dollar change in home values (between 2006 and 2009) for counties in a given quintile of financial distress as measured by our CL80 measure. The horizontal line represents the MPC estimated by Mian et al. [2013]. In general, our estimates are slightly smaller.²⁴ More importantly, the MPC increases with

²³One minor difference against the regressions of Mian et al. [2013] is that, where they calculate ΔHV_t^i and the change in financial wealth ΔFW_t^i by multiplying the initial value in each by the percentage change in corresponding market indices, we take the direct difference in zip-code level home values and financial wealth between 2006 and 2009. This affords us broader coverage, and is possible because we now have access to home value and IRS SOI data for 2009 that they did not. The full construction of these variables is described in Appendix A and particularly subsection A.2.

²⁴Related, Dupor et al. [2019] estimate that the MPC is 0.9 cents, which is smaller than the MPC estimated by Mian et al. [2013] but also in the range of our estimates.


Table 8: Auto spending at the zip-code level

FD Measurement taken in 2002:	Δ_{06-09} Auto Spending			
	(DQ30)	(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.

the incidence of FD from less than 1 cent to more than 2 cents.

Figure 15: Marginal Propensity to Consume out of a Dollar change in home prices by Quintile of DQ30 in 2002.



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 et al. \[2013\]](#) in their fifth column of Table 5.

7 Concluding Remarks

The main findings of this paper are that household-level FD is very unequally distributed, that FD affects consumer vulnerability to macroeconomic shocks, and that in some cases, FD and its dispersion matter for the aggregate consumer spending response to an economywide shock. Our paper proceeds by first establishing three facts: (i) regions in the U.S. vary significantly in their “FD-intensity,” measured either by how much additional credit households therein can access or in how delinquent they typically are on debts, (ii) shocks that are typically viewed as “aggregate” in nature hit geographic areas quite differently, and (iii) FD is an economic “pre-existing condition”: the share of an aggregate shock borne by a region is positively correlated with the level of FD present *prior* to the shock. Using an empirically disciplined and institutionally rich model of consumer debt and default, we show that in both the Great Recession and in the initial outcomes in the COVID-19 pandemic, FD mattered. Our model implies that the uneven distribution of FD implied uneven consumption responses. We find that this remains true even when breaking the positive correlation between initial FD and the incidence of an “aggregate” shock on each region.

In identifying FD as an amplifier of shocks—starting with household-level spending—our findings reinforce the message first discovered and conveyed by [Mian et al. \[2013\]](#) and [Mian and Sufi \[2010\]](#). Those authors were the first to show decisively that macroeconomic outcomes run through household balance sheets and credit health.

Our work suggests also that the state of households vis-à-vis their creditors, which we capture through FD, is also likely to be important in governing macroeconomic fragility in terms of aggregate consumer spending and provides information in addition to that encoded in leverage or net worth.

A conjecture for future work that emerges from our paper is that macroprudential policy may benefit from tracking either or both of the measures of FD we have provided. FD can be observed at a fairly granular level and hence may well be relevant to forecasting not only the severity of damage to local or regional consumption from macroeconomic shocks, but also the amplification of the shocks themselves.

References

- Aditya Aladangady. Housing wealth and consumption: Evidence from geographically-linked microdata. *American Economic Review*, 107(11):3415–46, 2017.
- Boragan Aruoba, Ronel Elul, and Sebnem Kalemli-Ozcan. How big is the wealth effect? Decomposing the response of consumption to house prices. 2018.
- Kartik Athreya, Juan M. Sánchez, Xuan S. Tam, and Eric R. Young. Labor market upheaval, default regulations, and consumer debt. *Review of Economic Dynamics*, 18(1):32–52, January 2015.
- Kartik Athreya, Juan M. Sánchez, Xuan S. Tam, and Eric R. Young. Bankruptcy and delinquency in a model of unsecured debt. *International Economic Review*, 59(2):593–623, 2017.
- Kartik Athreya, José Mustre-del Río, and Juan M. Sánchez. The persistence of financial distress. *The Review of Financial Studies*, 32(10):3851–3883, 2019.
- Adrien Auclert and Kurt Mitman. Consumer bankruptcy as aggregate demand management. 2019.
- David Berger, Veronica Guerrieri, Guido Lorenzoni, and Joseph Vavra. House Prices and Consumer Spending. *Review of Economic Studies*, 85(3):1502–1542, 2018.
- J. Y. Campbell and J. F. Cocco. How Do House Prices Affect Consumption? Evidence from Micro Data. *Journal of Monetary Economics*, 53(3):591621, 2007.
- Christopher Carroll, Jiri Slacalek, Kiichi Tokuoka, and Matthew N. White. The distribution of wealth and the marginal propensity to consume. *Quantitative Economics*, 8(3):977–1020, November 2017.
- Raj Chetty, John N. Friedman, Nathaniel Hendren, and Michael Stepner. How did covid-19 and stabilization policies affect spending and employment? a new real-time economic tracker based on private sector data, June 2020.

- Bill Dupor, Marios Karabarbounis, Marianna Kudlyak, and M. Saif Mehkari. Regional Consumption Responses and the Aggregate Fiscal Multiplier. Working Paper Series 2018-4, Federal Reserve Bank of San Francisco, February 2018.
- Bill Dupor, M. Saif Mehkari, Rong Li, and Yi-Chan Tsai. The 2008 US Auto Market Collapse. Technical report, 2019.
- Carlos Garriga and Aaron Hedlund. Mortgage Debt, Consumption, and Illiquid Housing Markets in the Great Recession. October 2017.
- Andrew Glover, Jonathan Heathcote, Dirk Krueger, and Jose-Victor Rios Rull. Controlling a pandemic. April 2020.
- Joseph Gruber and Robert Martin. Precautionary savings and the wealth distribution with illiquid durables. 2003.
- Fatih Guvenen, Serdar Ozkan, and Jae Song. The nature of countercyclical income risk. *Journal of Political Economy*, 122(3):621–660, June 2014.
- Juan Carlos Hatchondo, Leonardo Martinez, and Juan M. Sánchez. Mortgage defaults. *Journal of Monetary Economics*, 76:173–190, November 2015.
- Kyle F. Herkenhoff. The role of default (not bankruptcy) as unemployment insurance: New facts and theory. Working Paper, January 2013.
- Kyle F. Herkenhoff and Lee E. Ohanian. Foreclosure delay and u.s. unemployment. Working Paper, March 2012.
- Greg Kaplan and Giovanni L. Violante. How Much Consumption Insurance beyond Self-Insurance? *American Economic Journal: Macroeconomics*, 2(4): 53–87, October 2010.
- Greg Kaplan, Giovanni L. Violante, and Justin Weidner. The Wealthy Hand-to-Mouth. *Brookings Papers on Economic Activity*, 45(1 (Spring)):77–153, 2014.
- Greg Kaplan, Kurt Mitman, and Giovanni L. Violante. 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.

- Greg Kaplan, Benjamin Moll, and Giovanni L. Violante. Monetary Policy According to HANK. *American Economic Review*, 108(3):697–743, March 2018.
- Greg Kaplan, Kurt Mitman, and Giovanni L. Violante. The housing boom and bust: Model meets evidence. 2019.
- Greg Kaplan, Ben Moll, and Gianluca Violante. Pandemics according to hank. May 2020.
- P. Krusell and A Smith. Consumption-savings decisions with quasi-geometric discounting. *Econometrica*, 71(1):365–375, 2003.
- Moritz Kuhn and José-Victor Ríos-Rull. 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.
- Igor Livshits, James MacGee, and Michèle Tertilt. Consumer bankruptcy: A fresh start. *American Economic Review*, 97(1):402–18, March 2007.
- Atif Mian and Amir Sufi. The great recession: Lessons from microeconomic data. *American Economic Review*, 100(2):51–56, 2010.
- Atif Mian, Kamallesh Rao, and Amir Sufi. Household Balance Sheets, Consumption, and the Economic Slump. *The Quarterly Journal of Economics*, 128(4):1687–1726, 2013.
- Kurt Mitman. Macroeconomic Effects of Bankruptcy and Foreclosure Policies. *American Economic Review*, 106(8):2219–2255, August 2016.
- Christina Patterson. The matching multiplier and the amplification of recessions. 2018.
- Anthony Pennington-Cross. 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.
- Monika Piazzesi and Martin Schneider. Housing and macroeconomics. In John B. Taylor and Harald Uhlig, editors, *Handbook of Macroeconomics*, volume 2, pages 1547–1640. North-Holland Press, 2016.

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

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 what is shown in the paper. Table A1 shows some initial summary statistics for the entire data set.

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

Note: 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 over all the years considered in this study from Equifax presents a slight challenge: small random samples from FRBNY CCP/Equifax 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 could fix this issue, but the resulting datasets become difficult to process. Instead, then, we divide the zip codes for which we have IRS SOI data into 10 groups by population size 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.²⁵ In order to remain in our sample for a given quarter, individuals must be between 25 and 65 years old, inclusive.²⁶ 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 et al. [2013]. Net wealth is defined as the sum of housing wealth H and financial wealth FW less debt D . H is calculated 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.²⁷ This is done separately for zip codes and for counties. With a measure of total housing wealth in a geography thus defined, we calculate the housing leverage ratio as the total housing debt in a geography divided by the total housing wealth. Total housing debt is the mean housing debt, including both mortgages and home equity lines of credit²⁸ recorded in Equifax, in each geography multiplied by the number of households in that geography, taken from the Census.

To construct FW , we began by using IRS Summary of Income (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

²⁵Zip-code level data on CL80 and DQ30 are available at this [link](#) for the years 2006 and 2018.

²⁶Age is calculated using an individual’s recorded birth year, and so any records not including a birth year are also excluded.

²⁷To 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 et al. [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.

²⁸This includes both the home equity installment balance and the home equity revolving balance.

that fraction.²⁹ FW at the county level is simply calculated as the sum of FW in its component zip codes.³⁰ 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. Because 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. Next, we assign each geography a portion of the total debt from the SCF equal to that fraction.

In addition to the types of debt that Mian et al. [2013] tracks, we also include a measure of credit card debt at the zip-code and county levels. Here, we take the mean credit card debt by household in our Equifax sample and multiply that by the number of households in each geography.

As a measurement of consumption, 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 et al. [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 et al. [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.

By construction, then, the aggregate auto expenditures in our sample will accurately reflect the national difference from 2006 to 2009, but measurement error will be present at the local level to the extent that auto prices did not evolve in the same way across zip codes and counties from 2006 to 2009. If

²⁹Mian et al. [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.

³⁰To 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.

the price of pickup trucks dropped more than other types of cars, for example, and a particular rural county purchases mainly pickup trucks, then our data will underestimate the decrease in car consumption for that county just as [Mian et al. \[2013\]](#) did.

A.3 Financial Distress

Several measures of FD are used in this paper. 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.

With these two definitions in place, the remaining metrics used in the regression tables elsewhere make slight modifications to serve as robustness checks. “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 limit³¹ 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 any kind of debt recorded by the Federal Reserve Bank of New York/Equifax CCP.


Given that our sampling method over samples the smallest zip codes, we weight the aggregation of these four financial distress statistics to the county level by the number of households in each zip code.

³¹To give a clarifying example, say that there was an individual who in quarter 1 of 2006 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.

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 et al. \[2019\]](#). Thinking of a zip code as a collection of individuals, it follows that there should be some persistence in FD characteristics at a community level as well, although limited by the way that individuals sometimes move. In fact, however, FD at the zip-code level is more persistent than would be expected if individual-level persistence were the only factor at play. Figure A1 illustrates this point. Conditional on a zip code having been in the worst quintile of FD in 2000, 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, and occurs despite the fact that some or all of the original households that inhabited each zip code in 2000 could have moved out.

Figure A1: Persistence of Financial Distress at the Zip-code Level



Note: This graph is weighted in each year by the number of households in each zip code. “Random Chance, Quintile 5” presents what the probability of being in quintile 5 would be in each year if FD occurred randomly across zip codes.

Indeed, given from 2007 ACS data that the average person in the United


States will move about 12 times in their lifetime, and assuming those moves are distributed randomly over a 80-year lifetime, a back of the envelope calculation suggests that the average person in our sample moved 2.6 times in the years 2001-2017. If FD were only persistent at the individual level, then, and people had no tendency to sort themselves into zip codes with similar FD patterns, we would again expect the odds that a zip code in the worst quintile of FD in 2000 remains there in 2018 to be near random. Given that it is not, this shows evidence for additional mechanisms driving the persistence of zip-code level FD than simple persistence at the individual level. 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. In the case of the Great Recession, as shown in Figure A2, the share of debt in delinquency rose in tandem with the fall in home prices as people sought to smooth consumption by missing debt payments. Communities already in greater FD had less ability to use this channel for smoothing consumption. For each shock we consider, distinguishing zip codes that temporarily entered FD 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 for 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 Differences between FD and leverage


It is important to note that FD is not merely housing leverage repackaged, as may be wondered given some analogous findings of Mian et al. [2013]. Indeed, as shown in Figure A3, there does not appear to be a clear relationship between the two contemporaneously 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., more regions with more financial distress have lower leverage!

Figure A2: Evolution of Aggregate House Prices and Financial Distress



Note: The shaded area represents the Great Recession.

Figure A3: Correlation of Housing Leverage with FD (DQ30) in 2002



Note: 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.


A.3.3 Correlation between FD and the Housing Wealth Shock

In considering the implications of this drop in house prices for household balance sheets, it is useful to convey lost housing wealth as a fraction of net wealth. We follow [Mian et al. \[2013\]](#) in defining net wealth NW as the sum of housing wealth

H and financial wealth FW less debt D . In their framework, the housing net worth shock for a zip code i is then defined as the change in housing wealth $\Delta H_{i,06-09}$ divided by the initial net wealth NW_{06}^i .

Figure A4 documents the major fact to be established in this section: the incidence of the housing wealth shock upon zip codes was highly positively correlated with household FD. That is, higher FD in 2002 was associated with larger declines in housing wealth shocks in the Great Recession.

Figure A4: 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.

Notice that the housing-net-wealth shock decomposes into two contributing effects—the change in home prices and the share of wealth that was held in housing in 2006—as follows:

$$\frac{\Delta H_{i,06-09}}{NW_{06}^i} = \underbrace{\left(\frac{\Delta H_{i,06-09}}{H_{06}^i} \right)}_{\% \text{ chg. in house prices}} \underbrace{\left(\frac{H_{06}^i}{NW_{06}^i} \right)}_{\text{share of wealth in housing}}.$$

In terms of the change in home prices, Figure 6 shows that regions with more FD experienced greater percentage losses. These regions also tend to hold a larger share of their net wealth in their homes,³² so that they experienced a more severe

³²This can be seen, for example, in table 1. The difference between “Net Wealth Per HH” and “Net Fin. Wealth Per HH” is the housing wealth per household. Dividing this by the net wealth per household shows that housing wealth as a percentage of net wealth is monotonically increasing in FD: Quintiles 1,2,3,4, and 5 have housing shares of net wealth respectively equal to 37.5 percent, 40.7, 42.2, 42.8, and 52.2 percent.

housing net wealth shock along the second channel as well.

In other words, even if every zip code had experienced the same percentage decline in home prices, those with high FD would have tended to experience larger net wealth shocks. Similarly, this relationship would have held even if the share of wealth in housing had been held constant and only the home prices had been allowed to vary. As it is, the correlation at the zip code level between FD and the local housing net worth shock is highly robust because both of these effects are present and working in the same direction.

For example, the correlation is robust to alternative definitions of FD, as can be seen in Figure A6a. 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 A6b shows that the same relationship holds when measuring FD just before the recession started in 2006.

A.3.4 Amplification

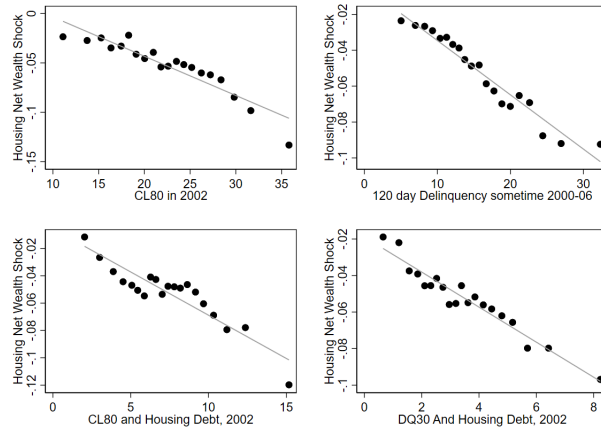
Changes in house prices are amplified differently into housing net worth shocks because of differences in leverage. A net worth shock is defined as in Mian et al. [2013] by

$$\frac{\Delta \log(p_{06-09}^{H,i}) H_{06}^i}{NW_{06}^i} = \underbrace{\left(\frac{\Delta \log(p_{06-09}^{H,i}) H_{06}^i}{H_{06}^i} \right)}_{\text{chg. in house prices}} \underbrace{\left(\frac{H_{06}^i}{NW_{06}^i} \right)}_{\text{share of wealth in housing}}.$$

Thus, the amplification is driven by two underlying components of the housing net worth shock, both of which themselves correlated with FD: the change in home prices and the share of wealth that was held in housing in 2006. Setting each component in turn constant to isolate variation in the other, it is possible to uncover the relative importance of each component to the overall housing net worth shock. Figure A7 plots the resulting relationship and shows that the effects of each are meaningfully correlated with FD. In other words, the observed relationship between housing price shocks during the Great Recession and FD

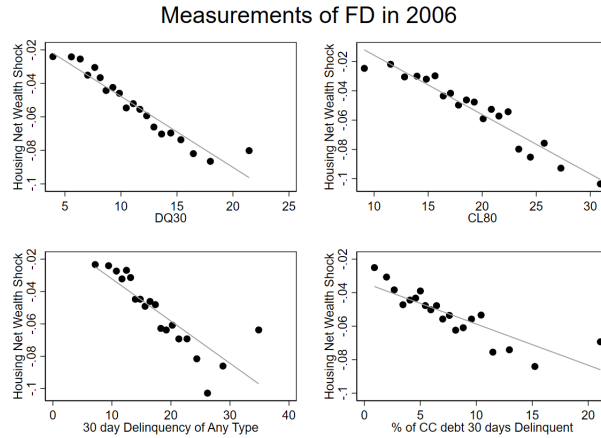
Figure A5: Robustness of the correlation between housing wealth shocks and FD

(a)



Note: “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.

(b)




Note: “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.

would have existed regardless of whether changes in home prices or the share of wealth people held in their homes were held fixed across the country.

Figure A7: Decomposition of 2006-09 House Price Shock



Sources: IRS SOI, Zillow, FRBNY Consumer Credit Panel/Equifax, Census Bureau. Quintile means of each object are weighted by that objects denominator in 2006; i.e., the “Housing Net Worth Shock” and “Shock Holding the Change in Home Price Constant” are weighted by net wealth, and the “Shock Holding the Share of Wealth in Housing Constant” is weighted by housing wealth.

B Recursive formulation of model

B.1 Nonhomeowner

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

$$N_{j,n}(a, z, \epsilon) = \max_{I_{rent} \in \{0,1\}} \left\{ I_{rent} R_{j,n}(a, z, \epsilon) + (1 - I_{rent}) B_{j,n}(a + e_n(z, \epsilon), z) \right\}, \quad (2)$$

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:

$$R_{j,n}(a, z, \epsilon) = \max \left\{ R_{j,n}^P(a, z, \epsilon), R_{j,n}^{BK}(a, z, \epsilon), R_{j,n}^{DQ}(a, z, \epsilon) \right\}, \quad (3)$$

and

$$B_{j,n}(a, z, \epsilon) = B_{j,n}^P(a, z, \epsilon). \quad (4)$$

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' :

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

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

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

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:

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

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

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

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:

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

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

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

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 $h' \in \{h_1, \dots, h_m\}$,

$$\hat{B}_{j,n}(a, z; h') = \max_{a', m'} u(c, h') + \beta_j E \left[H_{j,n-1}(h', m', a', z', \epsilon') | z \right] \quad (8)$$

$$\begin{aligned} s.t. \quad c + q_{j,n}^a(h', m', a', z)a' = \\ a + q_{j,n}^m(h', m', a', z)m' - I_{m' > 0} \xi_M - (1 + \xi_B)ph', \end{aligned}$$

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

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

$$B_{j,n}(a, z) = \max_{h' \in \{h_1, \dots, h_H\}} \hat{B}_{j,n}(a, z; h'). \quad (9)$$

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

B.2 Homeowner

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

$$H_{j,n}(h, m, a, z, \epsilon) = \max \left\{ P_{j,n}(\cdot), F_{j,n}(\cdot), D_{j,n}(\cdot), S_{j,n}^B(\cdot), S_{j,n}^R(\cdot) \right\} \quad (10)$$

where:

$$P_{j,n}(h, m, a, z, \epsilon) = \max \left\{ P_{j,n}^P(\cdot), P_{j,n}^{BK}(\cdot), P_{j,n}^{DQ}(\cdot) \right\}, \quad (11)$$

$$F_{j,n}(h, m, a, z, \epsilon) = F_{j,n}^P(\cdot), \quad (12)$$

$$D_{j,n}(h, m, a, z, \epsilon) = \max \left\{ D_{j,n}^P(\cdot), D_{j,n}^{BK}(\cdot), D_{j,n}^{DQ}(\cdot) \right\}, \quad (13)$$

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

$$S_{j,n}^R(h, m, a, z, \epsilon) = S_n^{R,P}(\cdot). \quad (15)$$

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:

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

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

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

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

problem:

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

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

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

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

$$P_{j,n}^{DQ}(h, m, a, z, \epsilon) = u(c, h) + \beta_j E \left[\gamma H_{j,n-1}(h', m(1 - \delta), 0, z', \epsilon') + (1 - \gamma) H_{j,n-1}(h', m(1 - \delta), a(1 + r^R), z', \epsilon') | z \right] \quad (18)$$

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

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

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:

$$F_{j,n}^P(h, m, a, z, \epsilon) = \hat{B}_{j,n}(a + ph(1 + \xi_B) - q_n^* m + e_n(z, \epsilon), z; h) \quad (19)$$

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:

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

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:

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

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:

$$D_{j,n}^{DQ}(h, m, a, z, \epsilon) = R_{j,n}^{DQ}(a, z, \epsilon) - \Phi. \quad (22)$$

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,

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

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,

$$S_{j,n}^{P,B}(h, m, a, z, \epsilon) = B_{j,n}(a + ph(1 - \xi_S) - q_n^* m + e_n(z, \epsilon), z). \quad (24)$$

B.3 Mortgage prices

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

$$q_{j,n}^m(h', m', a', z) = \frac{q_{pay,j,n}^m + q_{prepay,j,n}^m + q_{default,j,n}^m}{1 + r}. \quad (25)$$

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

$$q_{pay,j,n}^m(h', b', a', z) = \rho_n E \left[\text{mort pay, no def} + \text{mort pay, BK} + \text{mort pay, DQ} \middle| z \right], \quad (26)$$

with:

$$\begin{aligned} \text{mort pay, no def} &= I_{P_{j,n-1}^P}(h', m', a', z', \epsilon') \left[1 + (1 - \delta) q_{j,n-1}^m(h', m'', a'', z') \right], \\ a'' &= \hat{a}_{j,n-1}^{P,P}(h', m', a', z', \epsilon'), \end{aligned} \quad (27)$$

$$\text{mort pay, BK} = I_{P_{j,n-1}^{BK}}(h', m', a', z', \epsilon') \left[1 + (1 - \delta) q_{n-1}^m(h', m'', 0, z') \right], \quad (28)$$

and

$$\begin{aligned} \text{mort pay, DQ} = I_{P_{j,n-1}^{DQ}}(h', m', a', z', \epsilon') & \left[1 + (1 - \delta) \times \right. \\ & \left. \left(\gamma q_{j,n-1}^m(h', m'', 0, z') + (1 - \gamma) q_{j,n-1}^m(h', m'', a'', z') \right) \right], \end{aligned} \quad (29)$$

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

Here, ρ_n is the age-specific survival probability and 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^* :

$$\begin{aligned} q_{prepay,j,n}^m(h', m', a', z) = & \\ & E \left[\left(I_{F_{j,n-1}}(h', m', a', z', \epsilon') \right. \right. \\ & \left. \left. + I_{S_{j,n-1}^R}(h', m', a', z', \epsilon') + I_{S_{j,n-1}^B}(h', m', a', z', \epsilon') \right) q_{j,n-1}^* \middle| z \right]. \end{aligned} \quad (30)$$

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

$$\begin{aligned} q_{default,j,n}^m(h', m', a', z) = & \\ & \rho_n E \left[\frac{\left(I_{D_{j,n-1}}(h', m', a', z', \epsilon') \right) ph'(1 - \bar{\xi}_S)}{m'} \middle| z \right]. \end{aligned} \quad (31)$$

B.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:

$$q_{j,n}^a(h', m', a', z) = \frac{q_{pay,j,n}^a + q_{DQ,j,n}^a}{1 + r}. \quad (32)$$

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,

$$q_{pay,j,n}^a(h', m', a', z) = \rho_n E \left[\begin{aligned} & I_{R_{j,n-1}}^P(a', z', \epsilon') + I_{B_{n-1}}(a' + e_{n-1}(z', \epsilon'), z', \epsilon') \\ & + I_{P_{j,n-1}}^P(h', m', a', z', \epsilon') + I_{F_{j,n-1}}^P(h', m', a', z', \epsilon') \\ & + I_{D_{j,n-1}}^P(h', m', a', z', \epsilon') \\ & + I_{S_{j,n-1}}^{R,P}(h', m', a', z', \epsilon') + I_{S_{j,n-1}}^{B,P}(h', m', a', z', \epsilon') \Big| z \end{aligned} \right]. \quad (33)$$

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,

$$\begin{aligned}
q_{DQ,j,n}^a(h', m', a', z) = (1 - \gamma)(1 + r^R)\rho_n E \Bigg[& I_{R_{j,n-1}^{DQ}(a', z', \epsilon')} \times q_{j,n-1}^a(h_R, 0, a'', z') \\
& + I_{D_{j,n-1}^{DQ}(h', m', a', z', \epsilon')} \times q_{j,n-1}^a(h_R, 0, a'', z') \\
& + I_{P_{n-1}^{DQ}(h', b', a', z', \epsilon')} \times q_{j,n-1}^a(h', m'', a'', z') | z \Bigg]
\end{aligned}
\tag{34}$$

with: $a'' = (1 + r^R)a'$ 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 $h' > h_R$.

C COVID-19 related income losses

How do we pin-down COVID-19 related income losses? There are three steps in our strategy. First, we compute the share of workers in “leisure and hospitality” for each quintile. This share corresponds to the “most affected” workers. We use the Census LEHD Origin-Destination Employment Statistics (LODES) data described in Figure 5 together the classification of zip-codes to quintiles of FD obtained from FRBNY Consumer Credit Panel/Equifax. Table A2 shows that while the least distressed areas have a leisure and hospitality employment share of around 9 percent, the most distressed areas have “most affected” employment shares above 12 percent. Thus, the U.S. average masks substantial heterogeneity by FD quintiles.

To obtain the decline in income for this group, we use the data in Figure 4. During the months of April and May, the decline of this sector was more than 40 percent. By June, there was a recovery but the seasonally adjusted decline was still larger than 30 percent. We use 30 percent as a conservative estimate for the rest of the year, because starting in July there was a new wave of cases and some states reinstated some restrictions, especially in the leisure and hospitality sector.

The second step is to identify workers not affected by the lockdown. To

Table A2: Employment Shares in “leisure and hospitality”

Quintile of FD	Employment share Leisure and Hospitality
1	9.2
2	10.0
3	10.6
4	11.1
5	12.1

Source: Census LODES and FRBNY Consumer Credit Panel/Equifax.

calculate the share of workers not affected by quintile of FD, we leverage data from the 10th wave of the Census Pulse Survey along with our FD measures from Equifax (used in Figure 8). First, using the Pulse Survey, we calculate state-level shares of individuals who report “no earnings losses since March 13, 2020 (for self or household member).” Second, we merge these state-level responses to our preferred Equifax FD measure (DQ30). Third, we estimate a simple linear regression relating state-level responses to state-level FD (weighted by state population). Finally, we use the estimated coefficients of this regression to impute quintile-specific shares of those not affected using each quintile’s average level of FD. The results of the regression appear in Table A3, while the imputed share of not affected workers by quintile of FD appear in the second column of Table 6.

Table A3: State-level percentage unaffected vs state-level FD

Variable	Coefficient
FD	-0.92 (0.42)
constant	60.07 (4.93)
N	51
R^2	0.08

Note: Robust standard errors appear in parentheses.

The last step is to derive how affected households are that belong neither to the most affected nor to the unaffected sectors. The reasoning here is very simple. We want to select a decline in income such that the aggregate decline in the economy is consistent with the predicted decline of 7.0 percent by the

consulting firm Macroeconomic Advisers. Notice that in this we need to take into account differences in average income across quintiles. We need a decline in income for the “somewhat affected” households of 9.1 percent to finalize our calibration.

D Regressions controlling by leverage

Next, we show how our main regression changes when controlling by leverage. The crucial point is that the estimates of the interaction terms are very similar to our benchmark results.

Table A4: 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* (0.01)	-0.013* (0.01)	-0.015* (0.01)	-0.013* (0.01)
FD	-5.458 (3.13)	-7.239* (2.89)	-7.495* (3.32)	-4.548* (2.02)
Δ_{06-09} Home Value \times FD	0.104*** (0.02)	0.068*** (0.02)	0.097*** (0.02)	0.073*** (0.01)
Housing Leverage Ratio ₀₆	-0.228 (1.15)	-1.216 (1.69)	-0.953 (1.48)	-0.677 (1.18)
Δ_{06-09} Home Value \times Housing Leverage Ratio ₀₆	0.018* (0.01)	0.014 (0.01)	0.016* (0.01)	0.019* (0.01)
Housing Leverage Ratio, 2006 \times FD	-0.320 (6.69)	4.519 (6.16)	4.164 (7.11)	1.637 (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.

E County-Level IV regressions

To mitigate the risk that their results stem from an omitted variable correlated with the decline in home prices, [Mian et al. \[2013\]](#) 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 [A5](#) and [A6](#), where we present the first and second stages of the regression as we do in Table [8](#) but instead at the country level

Table A5: First-Stage Regression, County-level data

FD Measurement taken in 2002:	Δ_{06-09} Home Value			
	(DQ30)	(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 A6: Second Stage regression, County-Level

FD Measurement taken in 2002:	Δ_{06-09} Auto Spending			
	(DQ30)	(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

Note: 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.

It may be worrisome that there is another variable correlated with our measures of FD that better summarizes a households' financial condition. The housing leverage ratio in particular is frequently suggested as a possible source of error, so Figure 15 directly compares our baseline to the results controlling for leverage and Table A4 shows the corresponding regression output.

Overall, these empirical results support the quantitative mechanisms highlighted in the previous subsections. Moreover, they are also consistent with the recent literature on consumption responses to house price shocks as exemplified by Mian et al. [2013] and Aladangady [2017], among others. However, these results are not intended to establish a causal relationship between financial distress and observed consumption declines. Indeed, our model suggests financial distress is a useful summary statistic capturing a history of high borrowing costs induced, in part, by impatience. Rather, these results corroborate our model's quantitative implications.