

Why Do Borrowers Default on Mortgages? A New Method For Causal Attribution

Peter Ganong and Pascal Noel*

June 24, 2021

Abstract

There are three prevailing theories of mortgage default: strategic default (driven by negative equity), cash-flow default (driven by negative life events), and double-trigger default (where both negative triggers are necessary). It has been difficult to test between these theories in part because negative life events are measured with error. We develop a new method for addressing this measurement error using a comparison group of borrowers with no strategic default motive. Our central finding is that only 6 percent of underwater defaults are caused exclusively by negative equity, about an order of magnitude lower than previously thought. We then analyze the remaining defaults. We find that 70 percent are driven *solely* by negative life events (i.e., cash-flow defaults), while 24 percent are driven by the *interaction* between negative life events and negative equity (i.e., double-trigger defaults). Together, these results provide a full decomposition of the three theories underlying borrower default and suggest that negative life events play a central role.

*ganong@uchicago.edu, pascal.noel@chicagobooth.edu. We thank John Campbell, Joao Cocco, and Marco Giacomelli for generously sharing code and for very helpful comments. We further thank Joao Cocco, Amir Kermani, and Paul Willen for serving as discussants on this paper. We also thank Kanav Bhagat, Neil Bhutta, Adam Guren, Kyle Herkenhoff, Peter Hull, Erik Hurst, Koichiro Ito, Anil Kashyap, Ben Keys, David Matsa, Neale Mahoney, Atif Mian, Jack Mountjoy, Mikkel Plagborg-Moller, Matthew Notowidigdo, Christopher Palmer, Jesse Shapiro, Amir Sufi, Joseph Vavra, Robert Vishny, Paul Willen, Luigi Zingales, and Eric Zwick for helpful conversations. We thank seminar participants at AEA, AREUEA, the Becker Friedman Institute, Berkeley Haas, Brown, BYU, the Central Bank of Ireland, Copenhagen University, CUNY, Dartmouth, Duke, Federal Reserve Bank of Boston, FHFA, GBURES, HBS, HEC, LBS, MFA, Michigan, Minnesota, MIT Sloan, NBER Summer Institute, NYU, the Stanford Institute for Theoretical Economics, TAMU, University of Chicago Law School, and UCLA for helpful comments. We thank Ari Anisfeld, Rei Bertoldi, Therese Bonomo, Guillermo Carranza Jordan, Lei Ma, Roshan Mahanth, Peter Robertson, and Nicolas Wuthenow for excellent research assistance. This research was made possible by a data-use agreement between the authors and the JPMorgan Chase Institute (JPMCI), which has created de-identified data assets that are selectively available to be used for academic research. All statistics from JPMCI, including medians, reflect cells with at least 10 observations. The opinions expressed are those of the authors alone and do not represent the views of JPMorgan Chase & Co. While working on this paper, the authors were compensated for providing research advice on public reports produced by the JPMCI research team. We gratefully acknowledge the financial support of the Center for Research in Security Prices, the Fama-Miller Center, the Initiative on Global Markets, and the Fujimori/Mou Faculty Research Fund at the University of Chicago Booth School of Business.

1 Introduction

“To determine the appropriate public- and private-sector responses to the rise in mortgage delinquencies and foreclosures, we need to better understand the sources of this phenomenon. In good times and bad, a mortgage default can be triggered by a life event, such as the loss of a job, serious illness or injury, or divorce. However, another factor is now playing an increasing role in many markets: declines in home values.” (Ben Bernanke, May 2008)

Mortgage defaults soared during the Great Recession, precipitating the worst financial crisis since the Great Depression. As Ben Bernanke explained in the speech quoted above, a key challenge facing lenders and policymakers addressing this crisis was understanding *why* defaults soared. On the one hand, Bernanke notes that defaults could be triggered by a negative life event such as the loss of a job. Indeed, the Great Recession saw a dramatic deterioration in the labor market, with the highest long-term unemployment rates ever recorded. On the other hand, Bernanke worried that a second factor, negative equity, was playing an increasing role. Indeed, house prices also fell dramatically during the crisis, leaving one in four borrowers underwater. The ambiguity over which force was pivotal in driving borrower default decisions made it difficult to develop an appropriate response.

Although Bernanke went on to describe the challenge of distinguishing between negative equity and negative life events as “novel,” housing economists recognized it as just the latest chapter in a longstanding debate between three theories of mortgage default. The first theory, dating back to Foster and Van Order (1984), is that default is triggered exclusively by negative equity. Because default according to this theory is a function of the house’s asset value but not the borrower’s personal financial situation, it has sometimes been called “strategic” default. The second theory, dating back to Riddiough (1991), is that mortgage default is triggered exclusively by a negative life event. According to this “cash-flow” theory, borrowers default when a negative life event reduces their cash flows, making it difficult to afford the mortgage payment. The third theory, dating back to Goldberg and Capone (1998), is that defaults are driven by an interaction between negative equity and negative life events. According to this “double-trigger” theory, both triggers are necessary to cause defaults.

The goal of this paper is to help distinguish between these theories. Our central finding is that only 6 percent of underwater defaults are caused exclusively by negative equity, about an order of magnitude lower than previously thought. We then further decompose the remaining defaults. We find that 70 percent are driven *solely* by negative life events (i.e., cash-flow defaults), while the remaining 24 percent are driven by the *interaction* between negative life events and negative equity (i.e., double-trigger defaults).

This topic has been the subject of substantial prior research. Nevertheless, disentangling the contribution of negative life events from that of negative equity remains one of the “central questions in this literature” according to a recent review article (Foote and Willen 2018). The question has remained central in part because of two problems: data limitations and a measurement challenge.

First, mortgage servicing data do not contain information on current income or possible triggering life events. Prior research has typically used measures that are out of date (such as the household’s payment-to-income ratio at origination) or geographically coarse (such as the regional unemployment rate). These coarse indicators are imperfect measures of the life events actually faced by individual borrowers at the time of default. This data limitation may lead to attenuation bias in the estimated role that life events play in explaining default.¹ A lack of data has been an enduring challenge for the literature. A review article by Vandell (1995) says that to make progress explaining mortgage default, researchers should build a dataset with details about a household’s financial circumstances at the time of default. In their review article published over two decades later, Foote and Willen (2018) call for constructing almost exactly the same dataset.

Second, even with data on current income, it is unclear what qualifies as a negative life event that is sufficiently important so as to cause a borrower to default. Is any drop in income sufficient or must income drop by a specific amount (e.g, at least 10 percent)? Can the drop in income be short lived, or must it last a specific amount of time, such as at least three months?

We overcome these two challenges with new data and a new empirical method. First, to overcome the data limitation, we link mortgage servicing records with bank account records, both from the JPMorgan Chase Institute (JPMCI). Bank account income is a rich (albeit noisy) measure of a household’s financial circumstances. Using mortgage servicing data for the same households, we measure home equity and mortgage default. The linked dataset has 3.2 million borrowers, which is orders of magnitude larger than previously-used data on a household’s financial situation at the time of default.

Second, to overcome the measurement challenge, we use the evolution of income of *above water defaulters* as an empirical benchmark of a default caused by a life event. The intuition for our approach is that borrowers with positive home equity cannot possibly be defaulting because of negative equity. Instead, they must be defaulting because of a negative life event.²

¹A related challenge is measurement error of a home’s value, which would in turn lead to error in measuring a borrower’s home equity. However, it is more straightforward to adjust for such measurement error because of the availability of validation data that contains both estimated home values and actual home sales prices. We show that our findings are unchanged when we implement such an adjustment using a two-sample instrumental variables approach in Section 4.3.

²It may initially be surprising that any borrowers with positive equity ever default (which we define in

Consistent with this view, we show that their income declines sharply in the months leading up to default. This decline therefore provides a benchmark for what it looks like when we can be confident that a default is caused by a life event. We can then compare the evolution of income for underwater borrowers to this benchmark.

We document a surprising pattern: the decline in income leading up to default by underwater defaulters—whose reasons for default are not known a priori—is nearly indistinguishable from that of above water defaulters—whose defaults must be caused by a life event. This reduced-form evidence that underwater defaulters experience the same financial distress as above water defaulters is qualitatively consistent with a central role for negative life events in explaining default.

What causal statement, if any, can be made using this reduced-form evidence? We answer this question in two steps: a potential outcomes model and an identification method. First, we present a simple causal model that formalizes the three longstanding theories of default in terms of potential outcomes (Rubin 1974). The model is summarized in Table 1 (and described formally in Section 2). We assume that either negative equity or a negative life event is necessary for default. We label a default as (1) “strategic” when negative equity is a necessary and sufficient condition for the default; (2) “cash-flow” when a negative life event is necessary and sufficient; and finally, (3) “double-trigger” when both negative equity and a negative life event are necessary.³ The model provides a unified, empirically tractable framework for decomposing the share of defaults attributable to each of these three theories.

Table 1: Theories of Mortgage Default

Default theory	Potential outcomes interpretation	Prior estimates	Our findings
Strategic	Negative equity is necessary and sufficient	30-70%	6%
Cash-flow	Negative life event is necessary and sufficient	No prior estimates	70%
Double-trigger	Both negative equity and negative life event are necessary	30-70%	24%

Notes: This table reports estimates of the share of underwater default accounted for by each theory.

Second, we show that the evolution of income before default can be used to identify our main specification as missing three mortgage payments, following the prior literature). In a frictionless environment borrowers would be able to instantaneously sell their home or seamlessly tap into their home equity to avoid missing payments after experiencing a negative life event. However, matching frictions make it difficult to sell quickly, and institutional frictions make it difficult and sometimes impossible for those who are unemployed or liquidity-constrained to quickly access illiquid housing wealth (Boar, Gorea and Midrigan 2020; DeFusco and Mondragon 2020). As a result, missed mortgage payments are ubiquitous for borrowers with positive equity (Low 2018).

³The label of “strategic default” as one meant to maximize a borrower’s financial wealth, irrespective of any negative life event, goes back to Riddiough and Wyatt (1994). This original meaning is commonly used in recent literature (see e.g. Bhutta, Dokko and Shan 2017), and our formal definition follows in this tradition. However, we note that there are other interpretations of the term “strategic default.” We discuss how our results relate to several alternative definitions in Section 5.

the share of strategic defaults within this causal model. Specifically, we argue that under plausible assumptions, comparing the path of income for above and underwater borrowers can separate the strategic defaults in row 1 (where life events are irrelevant) from the cash-flow and double-trigger defaults in rows 2 and 3 (where life events are necessary).

Our identification method relies on two ingredients to overcome the challenge that life events are difficult to observe. The first ingredient is the change in bank account income, which we assume is a noisy measure of negative life events. The second ingredient is a group whose defaults we assume must have been caused by negative life events. This is above water defaulters.

We combine these two ingredients using Bayes rule. The standard approach to causal inference puts the outcome on the left-hand side and treatment on the right-hand side of a regression equation. However, as is well known, this approach suffers from attenuation bias when treatment is measured with error. In our application, if negative life events are imperfectly observed at the individual level, simply regressing default on a noisy measure of such events will lead researchers to underestimate their importance in driving defaults. We therefore use Bayes rule to move *treatment* to the *left-hand* side, where noise will result in larger standard errors but not attenuation bias. In contrast to the standard approach, which would require studying the outcome conditioning on (noisily measured) treatment, we instead study treatment conditioning on the (precisely measured) outcome. Our approach draws inspiration from a labor economics literature which uses reverse regression to address concerns about attenuation bias.

This use of Bayes rule has a simple interpretation when applied to mortgage default. We condition on default and put the measure of life events on the left-hand side, comparing the evolution of bank account income for various groups. At one extreme, if the income drop for underwater defaulters looked exactly like that of above water defaulters, who we know are defaulting due to a life event, then we would conclude that a life event is also a necessary condition for all underwater defaults. At the other extreme, if underwater defaulters had the same change in income as underwater non-defaulters, then we would classify all their defaults as strategic.

When we interpret our reduced-form empirical findings through the lens of this econometric method, our estimates show that at most 6 percent of underwater defaults are strategic. In other words, we find that negative life events are a necessary condition for 94 percent of mortgage defaults, so 94 percent of defaults must be either cash-flow or double-trigger. Although concern about borrowers walking away from their homes solely due to negative equity was widespread (e.g., see Roubini 2008), our point estimate and confidence interval show little evidence of this type of default. Indeed, our central estimate is likely conservative.

We show that alternative specifications yield estimates of strategic default prevalence of 3 to 4 percent.

Our finding of little strategic default holds for different quantiles of income, time periods, loan-to-value (LTV) cutoffs, and definitions of default, and is robust to measurement error in LTV. First, the result does not depend on our choice of the mean as a summary statistic; we find very little strategic default across the income distribution. Second, we see almost no strategic default in every year between 2008 and 2015. Third, this finding does not depend on choosing an LTV cutoff at exactly 100 percent. Indeed, our test detects little evidence of strategic default until borrowers have LTVs well above 200 percent. Fourth, our baseline specification follows the prior literature in defining default as three missed payments, but we observe similar patterns in the data when we define default as foreclosure initiation or as one, two, four, or five missed payments. Finally, our results change little after accounting for measurement error in LTV using a two-sample instrumental variables approach.

Next, we complete the decomposition of the role for different theories of default by separating the cash-flow defaults (row 2 in Table 1) from the double-trigger defaults (row 3 in Table 1). This requires one additional ingredient: the causal impact of negative equity on default. We estimate this causal impact by adopting methods from the literature based on long-run regional variation in house price cyclicity (Palmer 2015, Guren et al. 2021). Although some papers have raised the possibility of cash-flow defaults (e.g., Riddiough 1991), the consensus view in the literature has been that any underwater default that is not strategic must be double-trigger.⁴ And indeed, we do find evidence for substantial double-trigger behavior: both triggers are necessary for 24 percent of underwater defaults. However, we find that pure cash-flow defaults, driven *solely* by negative life events, are also important: they account for 70 percent of underwater defaults. While such defaults have always been understood to be theoretically possible, we know of no prior empirical estimates of their prevalence among underwater borrowers.

The decomposition in Table 1 allows us to compare negative equity and negative life events on an equal footing and reveals that negative life events are far more important than negative equity. Our estimates imply that eliminating negative equity would prevent only 30 percent of defaults (strategic plus double-trigger) while eliminating life events would prevent 94 percent of defaults (cash-flow plus double-trigger).⁵

Our most surprising finding may be the near-absence of strategic default, despite prior leading estimates that between 30 and 70 percent of Great Recession defaults were strategic

⁴See Foote and Willen (2018) for a review of the recent literature. This literature focuses on double-trigger default theories as the alternative to strategic defaults.

⁵These estimates complement a literature which analyzes the effect of variation in penalties for debt non-repayment on bankruptcy (Dávila 2016; Gross et al. 2019; Indarte 2019).

(Gerardi et al. 2018; Guiso, Sapienza and Zingales 2013; Bhutta, Dokko and Shan 2017). To help explain *why* our estimates are so much lower, we re-analyze publicly available survey data from the Panel Study of Income Dynamics (PSID). Although the administrative bank data is key for obtaining precise estimates, we show that the survey data also yields similar conclusions about the lack of strategic default. Furthermore, this exercise allows us to compare our estimates to prior methodologies and prior definitions of strategic default. We demonstrate that using a comparison group to address measurement error (and not our data source or definition of strategic default) leads us to find less strategic default than prior work. Without an approach to address measurement error, attenuation bias leads to an underestimate of the importance of negative life events in explaining borrower default.⁶

Our finding of almost no strategic default may also be surprising because existing structural models predict substantial strategic default by deeply underwater borrowers. We use a benchmark structural model of mortgage default first developed in Campbell and Cocco (2015) to illustrate this point. We find that the model’s predictions closely match the data for borrowers with LTV less than 120. However, as LTV rises above this threshold, the model predicts that borrowers will default even in the absence of income shocks, which contrasts sharply with our empirical findings.

Despite this divergence between model and data for deeply underwater borrowers, we show that a simple extension can reconcile the two. Specifically, this type of model allows for the possibility that defaulting incurs a utility cost. However, Campbell and Cocco explain that the main difficulty with incorporating this cost is that there has previously been little data to discipline this parameter.⁷ We propose to use income losses before default as empirical targets. The reluctance of borrowers to default on a substantially underwater asset in the absence of income shocks is informative about how costly they perceive this default to be. We estimate that defaulting must impose a cost equivalent to a 25 percent decrease in the constant-equivalent consumption stream; we discuss potential sources of this cost in

⁶The literature analyzing regional unemployment rates and default exemplifies this attenuation bias. A long literature beginning with Campbell and Dietrich (1983) finds that regional unemployment has modest predictive power for default, which has been interpreted as consistent with a large role for strategic default. In the Great Recession, Goodman et al. (2010) used regional unemployment and titled their study: “Negative equity trumps unemployment in predicting mortgage default”. Yet as researchers have acquired better data and developed improved simulation methods, new evidence has emerged suggesting that unemployment may in fact be an important driver of mortgage default. First, Gyourko and Tracy (2014) show that prior estimates suffer from attenuation bias because *regional* unemployment is a poor measure of *individual* unemployment status. Second, Bayer, Ferreira and Ross (2016) show that improved regional measures of unemployment (e.g. by race) lead to a stronger correlation between unemployment and default. Third, Hsu, Matsa and Melzer (2018) show that mortgage default by the unemployed is highly responsive to cash-on-hand.

⁷Indeed, prior work has estimated a wide range of costs, from as low as a 1.5 percent decrease in the constant-equivalent consumption stream to as high as a 70 percent decrease (Kaplan, Mitman and Violante 2020; Hembre 2018; Schelkle 2018; Laufer 2018).

Section 6. Once this cost is incorporated, we find that the model is able to closely match the data. A high default cost thus provides one plausible microfoundation for the behavior we observe empirically.

This approach to reconciling model and data may provide a blueprint for a wide class of macro finance models where borrower default decisions play a central role. For example, models with endogenous borrower default have been used recently to inform questions about macro-prudential regulation, the origins of the 2008 financial crisis, bankruptcy and foreclosure policy, and optimal mortgage security design.⁸ These types of models must take a stand on what triggers borrower default. Our empirical results suggest that realistic models will feature negative life events such as cash-flow shocks to be a necessary condition for most default. We demonstrate that incorporating a large utility cost of defaulting is one specific way to achieve this. More generally, models with endogenous borrower default might seek to target large income drops before default, even for deeply indebted borrowers.

The paper proceeds as follows. Section 2 describes our econometric framework. Section 3 describes the two datasets with income, home equity, and default (administrative bank data and PSID survey data). Section 4 shows that life events are a necessary condition for almost all defaults using the administrative data, which is contrary to theories of strategic default. It then shows how we decompose the remaining cash-flow and double-trigger defaults. Section 5 replicates our finding of little strategic default using the PSID. Section 6 explores implications for modeling mortgage default. Section 7 concludes.

2 Econometric Framework

We begin by first defining our main causal object of interest: the fraction of an outcome for which a particular treatment is a necessary condition. The standard method for identifying this object suffers from attenuation bias when treatment—in our case negative life events—is measured with error. Second, we describe a new identification method that addresses these challenges. Third, we describe how we apply the method to measure the fraction of mortgage defaults for which negative life events are necessary. Finally, we describe a complementary object: the fraction of mortgage defaults for which negative equity is necessary.

2.1 Standard Approaches to Causal Attribution

In many social science applications, researchers seek to measure “how much of event Y is attributable to binary treatment T^* ?” One precise answer to this question is the change

⁸For example, see Corbae and Quintin (2015), Mitman (2016), Kaplan, Mitman and Violante (2020), Guren, Krishnamurthy and McQuade (2019), Campbell, Clara and Cocco (Forthcoming), Greenwald, Landvoigt and Van Nieuwerburgh (2021), Diamond and Landvoigt (2019), and Garriga and Hedlund (2020).

in outcome Y from eliminating treatment T^* , which can be written as

$$\alpha \equiv \frac{E(Y) - E(Y_0)}{E(Y)} \quad (1)$$

where Y_0 is the potential outcome function evaluated in the absence of treatment. By measuring the share of the outcome eliminated in the absence of treatment, this estimand captures the share for which treatment is a necessary condition. Pearl (1999) and Rosenbaum (2001) are the first two papers we are aware of that formally study this estimand. Yamamoto (2012) says this estimand answers a *causal attribution* question.

The causal attribution parameter, which we define in equation (1), is of central interest in a broad class of economic applications. Recent examples that estimate α include: “how much of the correlation in wealth between parents and their children is due to nature versus nurture?” (Black et al. 2019; Fagereng, Mogstad and Ronning 2021), “what fraction of bankruptcies are caused by medical expenses?” (Gross and Notowidigdo 2011; Himmelstein et al. 2005; Dobkin et al. 2018), “what fraction of house price movements during the Great Recession is explained by a shift in beliefs?” (Kaplan, Mitman and Violante 2020), and “what share of layoffs in the Great Recession can be accounted for by credit supply contractions?” (Chodorow-Reich 2014).

There are two standard methods for estimating α : counterfactuals in structural models and “back-of-the-envelope” calculations. The first method, as in the Kaplan, Mitman and Violante (2020) example in the prior paragraph, is to specify an economic environment and utility functions, estimate preferences and other parameters, and evaluate a counterfactual in a structural model using the estimated parameters. The second method relies on the observation that, because $Y = Y_1T^* + Y_0(1 - T^*)$, α can be rewritten as proportional to the product of the average treatment effect and the probability of treatment:

$$\frac{E(Y) - E(Y_0)}{E(Y)} = \frac{E(Y_1T^*) + E(Y_0(1 - T^*)) - E(Y_0)}{E(Y)} = \frac{E(Y_1 - Y_0|T^* = 1)P(T^*)}{E(Y)}. \quad (2)$$

This method is used in the other examples in the prior paragraph. In these other examples, researchers estimate a Local Average Treatment Effect, assume that it extrapolates to the population average treatment effect on the treated ($E(Y_1 - Y_0|T^* = 1)$), multiply by an estimate of the probability of treatment ($P(\hat{T}^*)$) and divide by an estimate of the average level of the outcome ($E(\hat{Y})$). This is sometimes called a “back-of-the-envelope” calculation. Applied to analyzing the fraction of mortgage default attributed to life events, if a researcher knows the average treatment effect of life events on the probability of default (e.g., $E(Y_1 - Y_0|T^* = 1) = 0.1$), the probability of life events (e.g., $P(T^*) = 0.3$) and the population

default rate (e.g., $E(Y) = 0.06$). they can use equation (2) to hypothetically conclude that 50 percent of default is attributable to life events. However, in settings where treatment is particularly difficult to observe, measurement error can bias estimates of α both because of attenuation bias in the estimate of treatment effects and because the probability of treatment may be unknown.

The twin challenges of extrapolation and measurement error in treatment also make it difficult to determine how much of mortgage default is causally attributed to life events. Recall Bernanke’s speech where he enumerated a series of possible life events and emphasized the importance of understanding their role in driving mortgage default. Why is this hard to do using the standard approach? First, although there are already causal estimates of the impact of *some* life events on default (Hsu, Matsa and Melzer 2018 studies unemployment, Gupta et al. 2017 studies cancer diagnosis), the causal impact of *all other* life events on default may be larger or smaller. Second, there is considerable uncertainty about the probability of a negative life event. Gerardi et al. (2015) estimates that among mortgagors (defaulters and non-defaulters), the probability of a negative life event ranges from 4 percent under a stringent definition of a large decline in income to 57 percent under a broad definition which includes several types of shocks.

Because of these challenges, a review article by Foote and Willen (2018) suggests that a successful research program to understand the role of life events requires linking data on *all* major negative life events to mortgage servicing records. This has been unachievable to date because it is hard to obtain data on all such events. Indeed, some events (e.g., involuntary hours reductions, chronic illness) might be impossible to measure in existing datasets.

2.2 A New Method for Causal Attribution

We propose an alternative method to estimate α . Our method requires two ingredients. The first ingredient is a noisy measure of treatment. We assume the noise is orthogonal to the treatment and the outcome, as in the classical errors-in-variables problem (Wooldridge 2010). We use the change in bank account income as a noisy measure of life events in our application. We describe why this measure is noisy in the data description in Section 3. The second ingredient, which is novel, is a group whose outcome $Y = 1$ must have been caused by treatment T^* . In our application, this is the above water defaulters, whose defaults we assume must be caused by a negative life event.

Environment and assumptions

We assume there exists a population distribution (T^*, G, Y, T) where T^* is treatment, G is group status, Y is the outcome, and T is a candidate noisy measure of T^* . The first three variables are binary. In our mortgage default application, $T^* = 1$ indicates a

life event, $G = 1$ is negative home equity, $Y = 1$ is the transition to mortgage default, and T is the change in bank account income relative to one year ago. Assume potential outcome function $Y(T^*, G)$ (sometimes denoted as Y_{T^*G} for brevity), so each individual has four potential outcomes: $Y(0, 0)$, $Y(1, 0)$, $Y(0, 1)$, and $Y(1, 1)$.⁹ Additionally, assume there exists a potential outcome function $T(T^*, G, Y)$. The econometrician observes random draws from (G, Y, T) but T^* is unobserved.

Assumption 1 (default requires life event or negative equity): $Y(0, 0) = 0$ with probability one

In our application, this assumption says that the outcome of mortgage default requires either negative equity or a life event. This implies that a mortgage default *without* negative equity *must* have a life event ($P(T^* = 1 | Y = 1, G = 0) = 1$). The intuition is that a home with positive equity has financial value to the borrower, so defaulting is not in their long-term financial interest. Instead, it must reflect the impact of an adverse shock, which reduces the borrower’s ability to afford their mortgage payment. This assumption is standard in the mortgage default literature (e.g. Bhutta, Dokko and Shan 2017). Assumption 1 serves as a substitute for validation data in enabling us to identify the relationship between observable T and unobservable T^* .

An alternative way to think of this assumption is that, following the terminology from Bernanke’s speech, we define a “life event” ($T^* = 1$) as anything that causes an above water borrower to default. Applied to underwater borrowers, the potential outcomes model therefore quantifies the share of defaults that are *not* caused exclusively by negative equity. Thus, it may be useful to think of “life events” as a shorthand for all the non-negative equity causes of default.¹⁰

It may initially be surprising to some readers that above water borrowers ever default (which we define in our main specification as missing three mortgage payments, following the prior literature).¹¹ After all, if houses were a completely liquid asset, then above water default would never occur because a homeowner could borrow against their home to cover the mortgage payment.

⁹Relative to Section 2.1, which uses potential outcome function $Y(T^*)$, we now enrich the potential outcome function to have two arguments ($Y(T^*, G)$).

¹⁰One concrete example not included in Bernanke’s list is defaulting to become eligible for a mortgage modification. One estimate of this motivation comes from Mayer et al. (2014), which studies an episode where one mortgage servicer started offering generous mortgage modifications to borrowers in default. Analyzing the subset of borrowers who miss consecutive payments, the paper finds a one-time 10 percent increase in new defaults. This means that, in the quarter that the program was implemented, 10 percent of defaults for borrowers missing consecutive payments were motivated by a desire to get a mortgage modification.

¹¹In Section 4.3 we show that our results are robust to using alternative missed payment thresholds and to defining default as the date of foreclosure initiation.

Yet in practice, there are two frictions that make above water default quite common. First, there are substantial frictions to accessing home equity for borrowers in financial distress (Boar, Gorea and Midrigan 2020; DeFusco and Mondragon 2020). Underwriting for refinancing and second liens requires a good credit history and a documented “source of repayment” (Office of the Comptroller of the Currency 2005), which usually means proof of income or proof of substantial liquid assets (Fannie Mae 2011). In the Fannie Mae underwriting guide, unemployment insurance is not an acceptable source of income. An unemployed homeowner who needs a loan to cover her current mortgage payments would not meet the prevailing underwriting standard during our sample period. Second, borrowers may also choose to sell their home, but there are frictions in this process as well (Gilbukh and Goldsmith-Pinkham, 2021). For example, Guren (2018) documents that fewer than half of listed homes were sold within three months.

These two frictions mean that above water default is quite common. In Figure A-1, we corroborate Low (2018)’s finding that above water default is ubiquitous. The figure shows that even at the peak of the housing crisis, 40 percent of defaults were by above water borrowers. Furthermore, not only is above water default ubiquitous, but the economic risks from such default are substantial. Borrowers face an immediate credit score impact. The credit score decline from falling behind by three months on a mortgage is more than 80 percent as large as the decline from foreclosure and almost 60 percent as large as the decline from bankruptcy (Christie 2010). Perhaps the bigger risk, however, is foreclosure.

Above water foreclosures are common too; in fact, the rate of foreclosure starts among borrowers who have missed three payments is similar and high for above water and underwater borrowers. In our sample, we calculate that 40 percent of above water borrowers who fall behind by three months have a foreclosure initiation within one year (Table A-1). This is slightly below the foreclosure initiation rate of 45 percent for underwater borrowers.

Why would a lender foreclose when the value of the collateral exceeds the value of the loan? The high propensity to foreclose even on above water borrowers has both institutional and economic roots. Many mortgage servicers are not even allowed to consider home equity in the foreclosure decision because of rules made by the government-sponsored enterprises (GSEs).¹² Although the lender would prefer that an above water defaulter sell their home and repay the loan in full, they cannot instantaneously force a sale; instead, the foreclosure process is the legal mechanism by which the lender attempts to trigger a sale. Further, Low (2018) shows that above water foreclosures are an equilibrium outcome in a quantitative model with matching frictions in the home sale market that make the time-to-sell and resale

¹²See Fannie Mae (2020) Chapter E-3.2 for Fannie Mae guidance and U.S. Department of Housing and Urban Development (2019) Chapter III-A-2 for FHA guidance.

value uncertain. The GSE practice of ignoring home equity sets the standard for the industry. Even for non-GSE-owned loans, servicers who do not follow the industry standard face increased litigation risk.

As a matter of economics, it is not obvious whether lenders *should* foreclose more quickly on above water or underwater homes whose mortgages are in default. On the one hand, the return to the lender from foreclosing on an above water home is higher because the lender will likely recoup the full balance outstanding on the loan (which is unlikely to occur for an underwater loan). On the other hand, an additional month of waiting for an above water borrower to sell their own home may yield a higher sale price than a resale through foreclosure and also enables the lender to avoid the upfront administrative costs associated with foreclosure.

Assumption 2 (conditional exogeneity): $\{Y(0, 1), Y(1, 0), Y(1, 1)\} \perp T^* | G$

This assumption says the treatment of a life event is orthogonal to the potential outcome $Y(T^*, G)$ conditional on group G . This assumption is standard in the measurement error literature (as in the CEV framework in Wooldridge 2010) and in the literature studying the motivations for default (Bhutta, Dokko, and Shan 2017, and Gerardi et al. 2018). To understand the economic content of this assumption, it is useful to note that many models assume a private utility cost of defaulting, which is sometimes called a “moral” or “stigma” cost (Kaplan, Mitman and Violante 2017; Hembre 2018; Schelkle 2018; Laufer 2018). In such models, Assumption 2 implies that the probability of a negative life event is orthogonal to the moral cost (which governs the probability of default conditional on various economic circumstances).

This assumption allows for three types of heterogeneity that are important in the mortgage default context. First, it allows for the possibility that underwater borrowers are more likely to have negative life events than above water borrowers ($P(T^* = 1 | G = 1) > P(T^* = 1 | G = 0)$), consistent with the findings in Bhutta, Dokko and Shan (2017).¹³ Second, it allows for heterogeneity in the causal impact of a life event on default, consistent with findings in Gerardi et al. (2018) that underwater borrowers are more sensitive to income shocks than above water borrowers ($E(Y_{11} - Y_{01}) > E(Y_{10} - Y_{00})$). Third, it does *not* require that G is exogenous with respect to potential outcomes; instead, the role of G is similar to a covariate that we condition on in the analysis.

Although we believe that this orthogonality assumption is plausible, it is natural to wonder how our conclusions would change if this assumption did not hold. In Appendix D,

¹³This is also consistent with evidence in Bernstein (2019) and Gopalan et al. (Forthcoming), which finds that borrowers with negative equity are more likely to suffer income declines because of constrained mobility and financial distress.

we relax Assumption 2 by allowing for correlated, latent heterogeneity in the probability of a life event and the probability of default in the context of a simulation. We show that, even in this case, the bias is small because studying transitions to default differences out much of the latent heterogeneity.

Assumption 3 (noisy measure of treatment): (a) $T(T^*, G, Y) = T(T^*)$ and $\{T(0), T(1)\} \perp (T^*, Y, G)$, and (b) $E(T(1)) \neq E(T(0))$

Assumption 3a says that the potential outcome function for T is orthogonal to the other variables in the model (such as life event T^* , housing wealth G and debt repayment Y). Intuitively, it says that T is a noisy measure of T^* . This assumption has the same economic content as the CEV assumption in Wooldridge (2010); however, it is not identical because the CEV assumption uses a continuous latent variable, while the latent variable T^* in our framework is binary.

Assumption 3b says that income T falls on average for a borrower with a life event T^* . This assumption is analogous to the assumption in the instrumental variables (IV) literature that the instrument affects the probability of treatment.

Assumption 3 implies that when a life event *does* occur, above and underwater borrowers have the same average decline in income. This allows us to use T of above water defaulters, who always have $T^* = 1$ by Assumption 1, to learn about $P(T^*)$ for underwater defaulters. We can directly test whether this holds using one life event that we do observe in our data: unemployment. Table A-2 shows that the income drop conditional on unemployment is indeed statistically indistinguishable for borrowers with positive and negative home equity, consistent with Assumption 3. We discuss evidence consistent with an additional testable implication of Assumption 3 in Section 4.3.

Assumption 4 (monotonicity): $Y(1, 1) \geq Y(1, 0)$, $Y(1, 1) \geq Y(0, 1)$

This assumption says that a borrower is weakly more likely to default when they have both a life event and negative equity than when they have only one of the two.

We use these four assumptions to interpret the default motivations of underwater borrowers (borrowers with $G = 1$). Specifically, we ask what fraction of underwater mortgage default can be causally attributed to life events? We call this $\alpha_{\text{life event}}$.

Proposition 1: Under the environment described above and assumptions 1, 2, 3, and 4,

$$\alpha_{\text{life event}} \equiv \frac{E(Y|G = 1) - E(Y_{01}|G = 1)}{E(Y|G = 1)} = \frac{E(T|Y = 1, G = 1) - E(T|G = 1)}{E(T|Y = 1, G = 0) - E(T|G = 1)}. \quad (3)$$

Proof: See Appendix C.

The key step in the proof relies on Bayes rule. The standard approach to causal attribution puts the outcome Y on the left-hand-side and treatment T^* on the right-hand side of a regression equation. However, this approach suffers from attenuation bias when treatment is measured with error (i.e., we observe T instead of T^*). We therefore use Bayes rule to move the measure of treatment T to the *left-hand* side, where noise will result in larger standard errors but not attenuation bias. Hence, equation (3) relies on comparing this noisy measure of treatment across groups. We provide additional intuition for the formula in the context of mortgage default in Section 2.3.1.

2.2.1 Benefits, Costs, and Comparison to Prior Literature

The benefits and costs of this approach can be framed in terms of a bias-variance trade-off. The benefit is that it provides unbiased estimates of the causal attribution parameter, even when treatment is measured with error. The cost is that because there is measurement error on the left hand side of the estimating equation, the variance of the estimate will be larger. We demonstrate this trade-off using a simulation in Appendix E. Nevertheless, when we apply this approach to a large administrative dataset in Section 4, it still yields economically precise findings.

The approach has two additional types of limitations worth noting. First, it identifies fewer parameters of interest than the standard approach to measurement error. Although it identifies the product of the treatment effect and the probability of treatment (thereby identifying the central object of interest for the class of causal attribution problems), it does not distinguish between the two underlying inputs. For example, a finding of $\alpha_{\text{life event}} = 1$ in our application could be consistent with either infrequent negative life events that have a large impact on default or frequent ones that have a small impact on default. Because the main problem is that life events themselves are difficult to measure, we cannot distinguish between these two. To identify the treatment effect, in addition to Assumptions 1-4, we would also need to know the probability of treatment.

The second limitation is that the method imposes additional restrictions on the data-generating process which are not needed for the standard approach. First, it requires that the outcome variable (default) be binary and measured without error. This requirement is essential because we use both group status and the outcome to infer treatment. Second, it requires that the treatment can be modeled as binary. The appropriateness of this assumption is context-specific. In the mortgage default context, Guiso, Sapienza and Zingales (2013), Bhutta, Dokko and Shan (2017) and Gerardi et al. (2018) have modeled treatment as binary. In addition, we show that it is possible to offer a similar interpretation of the data without the binary treatment assumption using a structural model of mortgage default in

Section 6.

Our approach has some parallels with reverse regression, with difference-in-differences (DiD), and with instrumental variables (IV). It is similar to reverse regression in that we correct for attenuation bias in a regression by moving the noisy variable from the right-hand side to the left-hand side of the regression equation.¹⁴ However, reverse regression is most commonly used to reject or fail-to-reject a specific null hypothesis (e.g., “is there wage discrimination?”), while our approach provides a quantitative, causal interpretation of the sample moments.

The use of a comparison group (Assumption 1) is in some ways analogous to the control group in a DiD design. In DiD, the researcher estimates an average treatment effect by comparing the *outcome* for a treatment group and a control group. The idea of the control group is to directly identify what would have happened to the outcome if *no one* was treated. In our method, the comparison group gives a counterfactual for what would have happened to the *noisy measure of treatment* if treatment was a necessary condition for the outcome. We then use Bayes rule and other assumptions to identify what would have happened to the outcome if no one was treated.

Finally, our approach is similar to a linear IV setup in that we rescale a relationship of interest using an auxiliary equation. In an IV analysis, an auxiliary equation that relates the endogenous variable to the instrument is used to rescale the relationship between the outcome and the instrument. In our analysis, an auxiliary equation that measures what T would look like if everyone received treatment is used to rescale and interpret the value of T for the group of interest.

¹⁴The classic example of reverse regression is a researcher who seeks to measure the whether an employer is discriminating against their female employees in setting wages. The ideal test is to regress wages on employee productivity and gender, and see if women receive lower wages conditional on productivity. This test is infeasible because productivity is unobserved, but noisy measures of productivity (credentials) are typically observed. Reverse regression puts the credentials on the left-hand side, wages on the right-hand side, and tests if women have better credentials conditional on wages. The idea is that the noise in credentials will average to zero and therefore enable an unbiased estimate of discrimination. Informed by the critique of Goldberger (1984), reverse regression lost favor in the labor economics literature. One of his objections was that reverse regression will be biased when wages are measured with error, as they almost always are in survey datasets. However, this should not be an issue for our application because we observe administrative data on mortgage default.

2.3 Application to Mortgage Default: Causal Attribution of Negative Life Events

2.3.1 Formula and intuition

In the mortgage default application, our method is easy to use and interpret. We can write equation (3) as:

$$\alpha_{\text{life event}} = \frac{\overbrace{E(\Delta Inc^{\text{UnderwaterDefaulter}})}^{\text{group of interest}} - E(\Delta Inc^{\text{AllUnderwater}})}{\underbrace{E(\Delta Inc^{\text{AboveWaterDefaulter}})}_{\text{benchmark of 1: all defaults caused by life events}} - \underbrace{E(\Delta Inc^{\text{AllUnderwater}})}_{\text{benchmark of 0: no defaults caused by life events}}}. \quad (4)$$

The formula relies on comparing the change in income for underwater defaulters (our group of interest) to two benchmarks. The first benchmark is to determine what it would look like in the data if $\alpha_{\text{life event}} = 1$, in other words, if a life event was a necessary condition for every default. This benchmark draws on the assumption that a life event is indeed a necessary condition for all above water defaults. Therefore, the $\alpha_{\text{life event}} = 1$ benchmark is the change in income for above water defaulters, i.e. $E(\Delta Inc^{\text{AboveWaterDefaulter}})$. If the income drop for underwater defaulters was the same as that for above water defaulters, who must be defaulting due to a life event, then our assumptions imply that a life event is also a necessary condition for every underwater default ($\alpha_{\text{life event}} = 1$).

The second benchmark is what would it look like in the data if defaults were driven exclusively by negative equity ($\alpha_{\text{life event}} = 0$)? In this scenario, life events would be irrelevant for default. If life events were indeed irrelevant for default, then the average value of the noisy measure of life events (the change in income) would be the same for defaulters and non-defaulters. Hence the $\alpha_{\text{life event}} = 0$ benchmark is the change in income for all underwater borrowers, including both defaulters and non-defaulters, i.e. $E(\Delta Inc^{\text{AllUnderwater}})$. This benchmark arises intuitively from Foster and Van Order’s (1984) classic description of strategic default:

A key point about model 1 [a model of negative equity-driven default] is that personal characteristics of the borrower (income, employment status, etc.) are irrelevant. This is a characteristic of most option models of default;

Because income is “irrelevant” for this decision, it is natural to expect defaulters motivated by negative equity to have the same change in income as non-defaulters. Thus, if we observe the same income drop for underwater defaulters and nondefaulters (i.e. $E(\Delta Inc^{\text{UnderwaterDefaulter}}) = E(\Delta Inc^{\text{AllUnderwater}})$), then our assumptions imply that negative life events play no role in the default decision ($\alpha_{\text{life event}} = 0$).

Finally, if $\Delta Inc^{UnderwaterDefaulter}$ is in between the two benchmarks, then the share of defaults causally attributed to life events is between 0 and 1.

2.3.2 Potential Outcomes Interpretation

Our framework also allows us to separate default behavior in terms of three potential outcome types which correspond to specific economic mechanisms in the context of mortgage default. Recall that mortgage default Y is a function of two binary variables, life event treatment T^* and equity status G , i.e $Y_{\text{life event, negative equity}}$. We separate defaulters into three types based on possible combinations of these potential outcomes.¹⁵

These types have natural labels in terms of the prior literature, which we show in Table 2.

1. First, there are “strategic” (ST) defaulters for whom negative equity is a necessary and sufficient condition, as in Foster and Van Order (1984) and Bhutta, Dokko and Shan (2017). These borrowers would default solely due to negative equity ($Y_{01} = Y_{11} = 1$) but would not default solely due to a life event ($Y_{10} = 0$).
2. Second, there are “cash-flow” (CF) defaulters for whom a life event is a necessary and sufficient condition, as in Riddiough (1991). These borrowers would default with just a life event ($Y_{10} = Y_{11} = 1$) but would not default solely due to negative equity ($Y_{01} = 0$).
3. Third, there are “double-trigger” (DT) defaulters for whom a life event and negative equity are both necessary conditions, as in Goldberg and Capone (1998), Foote, Gerardi and Willen (2008), and Gerardi et al. (2018). These borrowers would default in response to both a life event and negative equity ($Y_{11} = 1$), but would not default if only one trigger was present ($Y_{01} = Y_{10} = 0$).

The estimand $\alpha_{\text{life event}}$ has a clear interpretation relative to these potential outcome types. As defined in equation (3), $\alpha_{\text{life event}}$ measures the fraction with potential outcome $Y_{01} = 0$ (i.e., the fraction of underwater defaults that would be eliminated in the absence of life events). It

¹⁵The fact that there are three potential outcome types arises from Assumptions 1 and 4 and our focus on defaulters who are already underwater. In principle, with four binary potential outcomes for each household $\{Y_{00}, Y_{01}, Y_{10}, Y_{11}\}$, there are $2^4 = 16$ possible combinations. We narrow the set of potential outcomes with four steps (two assumptions and by omitting two types that are irrelevant for our research question). First, Assumption 1 is that default requires either a life event or negative equity, so $Y_{00} = 0$ for all borrowers, which leaves $2^3 = 8$ possible combinations for the remaining three binary potential outcomes. Second, Assumption 4 (monotonicity) rules out any combination where Y_{01} or Y_{10} are 1 (i.e., the borrower would default with only a life event or with only negative equity) but $Y_{11} = 0$ (i.e., the borrower wouldn’t default with both a life event and negative equity). Third, we omit the type who never defaults where all potential outcomes are zero regardless of life events or negative equity. Fourth, we omit the type for whom either a life event or negative equity is sufficient for default (i.e., $Y_{10} = Y_{01} = Y_{11} = 1$). Because we are estimating α specifically among defaulters who *do* have negative equity, and this type would default solely due to this negative equity, life events are effectively irrelevant to their default decision.

therefore quantifies the combined share of cash-flow and double-trigger defaults (CF + DT) among underwater defaulters. These are defaults where life events are a *necessary* condition, or, in the language of Yamamoto (2012), defaults that can be “causally attributed” to life events. The balance, $1 - \alpha_{\text{life event}}$, therefore quantifies the share of strategic defaults (ST), those with potential outcome $Y_{01} = 1$.

2.4 Causal Attribution of Negative Equity in Mortgage Default

A related estimand captures the fraction of underwater defaults for which negative equity is a necessary condition:

$$\alpha_{\text{negative equity}} \equiv \frac{E[Y(T^*, 1) - Y(T^*, 0)|G = 1]}{E(Y|G = 1)}, \quad (5)$$

where the numerator is the causal impact of negative equity on default for underwater borrowers and the denominator is the probability of default for all underwater borrowers.

Combined, an estimate of $\alpha_{\text{negative equity}}$ along with an estimate of $\alpha_{\text{life event}}$ enable a full partition between the three theories of default in the potential outcomes model. While our estimate of $\alpha_{\text{life event}}$ quantifies the share of underwater defaults that would be eliminated in the absence of negative life events (those with potential outcome $Y_{01} = 0$, capturing the cash-flow [CF] and double-trigger [DT] defaults), an estimate of $\alpha_{\text{negative equity}}$ quantifies the share of underwater defaults that would be eliminated in the absence of negative equity (those with potential outcome $Y_{10} = 0$, capturing the strategic [ST] and double-trigger [DT] defaults). Therefore, using the fact that $ST + CF + DT = 1$, the strategic share is $ST = 1 - \alpha_{\text{life event}}$, the cash-flow share is $CF = 1 - \alpha_{\text{negative equity}}$, and the double-trigger share is $DT = \alpha_{\text{life event}} + \alpha_{\text{negative equity}} - 1$.

This decomposition is useful for two reasons. First, by quantifying the share of defaults that would be eliminated in the absence of each trigger, it allows us to compare the role of negative life events and negative equity on an equal footing. Second, by identifying the share of double-trigger defaults, it allows us to quantify the importance of *interaction* effects between the two triggers.

While the new method we propose in Section 2.2 is not suitable for estimating $\alpha_{\text{negative equity}}$, this estimand can be identified using the standard approach from Section 2.1. The standard approach is feasible for $\alpha_{\text{negative equity}}$ because the “treatment” of negative equity is more reliably observed and instruments for negative equity are readily available based on the prior literature. The key assumption of the standard approach is that we can extrapolate from the Local Average Treatment Effect of eliminating negative equity induced by the instrument to the population impact ($E[Y(T^*, 1) - Y(T^*, 0)|G = 1]$). In Section 4.6 we report estimates

of $\alpha_{\text{negative equity}}$.

3 Data

3.1 Linked Income-Servicing Data

Our primary analysis uses a novel administrative dataset from Chase that links checking account records and mortgage servicing records. These records are linked and then de-identified by Chase.¹⁶ This linkage is possible because Chase is both a consumer bank that offers checking accounts and a mortgage servicer.

Income in the checking account data captures a household’s post-tax cash flow each month, which is useful for understanding how cash flows affect mortgage default. Income is measured with error in the checking account data, albeit for different reasons than in surveys or tax data. For example, if a household transfers money in from a retirement account or receives a transfer of funds from a relative, this may look like income from the perspective of the checking account. If a household has multiple checking accounts, we compute income as the sum of all account inflows across all these accounts. We also observe some individual components of income, such as unemployment benefits.

A second key variable—the LTV ratio—comes from the mortgage servicing data. LTV is the ratio of total mortgage debt to estimated home value. Total mortgage debt, including second liens not serviced by Chase beyond the first lien, is observed reliably in the JPMCI data.¹⁷ Estimated home value is constructed using the standard procedure of inflating purchase price by a local measure of house price changes from CoreLogic. One potential concern is measurement error in a home’s value, which would in turn lead to error in measuring borrower’s home equity. It is straightforward to adjust for measurement error in a home’s value because of the availability of validation data that contains both estimated home values and actual home sales prices. We show that our findings are unchanged when we implement such an adjustment using a two-sample instrumental variables approach in Section 4.3.

We study borrowers who have cumulatively fallen behind on their mortgage by three monthly payments in most of our analysis. This is also known as 90 days past due. This is a common threshold for a mortgage to be considered in default (Foote and Willen 2018; Bhutta, Dokko and Shan 2017). However, we examine other thresholds for default in our robustness analysis.

¹⁶See Farrell et al. (2017) and Farrell, Bhagat and Zhao (2018) for JPMCI research using this linked dataset.

¹⁷Nineteen percent of Chase-serviced first liens had second liens in 2011, which is similar to 15 percent for a benchmark sample of first liens linked to credit bureau data called Credit Risk Insight Servicing McDash.

Our analysis uses first lien mortgages serviced by the bank between January 2007 and August 2015 and household income measured using checking accounts. The linked dataset has 3.2 million mortgages. In our main analysis we analyze defaults with an open checking account from one year before default through the date of default. This analysis sample contains 85,654 above water defaults and 52,331 underwater defaults from 2008 through 2015 (see Appendix B for details on sample construction). We also analyze the evolution of income for the universe of underwater borrowers, both defaulters and non-defaulters (657,000 borrowers). Finally, for our analysis of the impact of negative equity on default for computational reasons we subset to a 10 percent random sample of the full linked dataset (above water and underwater, defaulters and non-defaulters), which includes 316,353 borrowers and 13.2 million borrower-months.

Mortgage borrowers at Chase have similar characteristics to mortgagors in other, more widely-used datasets (McDash and CRISM). We discuss these datasets in more detail below. Table A-1 shows that the Chase data are similar to McDash and CRISM in terms of the default rate, the share of borrowers who are investors, the share of all borrowers who are underwater, and the share of defaulters who are underwater.

3.2 Other datasets

To check that our results are not unique to the Chase sample, we supplement our analysis with other data sets. For our main analysis on the role of life events in causing default we conduct parallel analysis using the best available public use data on income and mortgage default, which is survey data from the Panel Study of Income Dynamics (PSID). The PSID records pre-tax income y and consumption c in the calendar year prior to the survey. A supplement administered in 2009, 2011, and 2013 records housing costs (mortgage, property taxes, and insurance) m , home equity, and default (measured as 60 days past due) at the time of the survey. We follow the sampling choices from Gerardi et al. (2018): we drop households with LTV $>$ 250 percent, and we require that the head of household is in the labor force and between the ages of 24-65. We discuss the strengths and weaknesses of the PSID data compared to the JPMCI data in Section 5.

For our analysis on the share of defaults causally attributable to negative equity, we conduct parallel analysis using Credit Risk Insight Servicing McDash (CRISM) data. McDash is a dataset of origination and servicing records which covers about 70 percent of outstanding mortgage balances during our sample period. CRISM is a subset of the McDash data that can be linked to individual credit bureau records (and therefore can be used to calculate LTV ratios), which covers about 50 percent of outstanding mortgage balances during this time period. Our CRISM analysis on the causal impact of negative equity uses a 1 percent

random sample of borrowers with first lien mortgages who appear in the data between 2008 and 2015. This sample includes 386,000 borrowers (see Appendix B.2 for details on sample construction).

4 Main Results

Our main research design compares the evolution of mean income around default by home equity. We find that income declines are virtually indistinguishable for above and underwater defaulters. Using the econometric framework in Section 2, this empirical result implies that almost all defaults are causally attributable to negative life events, so the share of strategic defaults is small. We then estimate the causal impact of negative equity on default in order to distinguish between the remaining cash-flow and double-trigger defaults.

4.1 Research Design for Strategic Defaults

We compare underwater defaulters' income to two benchmarks using equation (4). In this section we discuss each of the two benchmarks. Then in Section 4.2 we use all three objects to estimate of the share of defaults causally attributable to life events.

The green triangles in Figure 1 show the evolution of income for above water defaulters in the twelve months prior to mortgage default. The x-axis is months since three missed payments. The figure is similar to an event study in that it shows monthly data relative to an event. It differs in that we focus specifically on income data *prior* to default. This choice is motivated by this paper's focus on the *causes* of the event, whereas the traditional event study is usually interested in the *consequences* of the event. The dependent variable is the change in monthly income relative to the average mortgage payment due one year before default.¹⁸ The time-series of income is consistent with Assumption 1, which says that above water defaults are caused by life events. Income falls sharply prior to default. Figure A-3 shows that average mortgage payment due is not rising prior to default, so the defaults we study are not triggered by changes in payment due.¹⁹

Figure 1 also shows the evolution of income for all underwater borrowers (both defaulters and non-defaulters). This series provides the benchmark for what income would look like

¹⁸This normalization facilitates the interpretation of the point estimates in terms of number of mortgage payments due and comparison of our data to the model in Section 6. Figure A-2 shows similar patterns (indeed, even less evidence of strategic default) when normalizing by prior income rather than by payment due.

¹⁹However, it would not be a problem for our methodology if the defaults in our sample were triggered by changes in payments. Payment changes are a non-negative-equity channel that could cause default and are thus considered a "life event" in the context of our model. Furthermore, payment changes directly affect borrower ability to pay.

if all defaults were strategic. We construct this series by re-weighting average income by calendar month to match the realized distribution of default dates. It is easiest to explain this scenario by first imagining a hypothetical world where all defaults occurred in a single month $s_{default}$. In this scenario, we would construct the series using average income in calendar months $\{s_{default} - 12, s_{default} - 11, \dots, s_{default}\}$. In practice, $s_{default}$ varies across borrowers. Let w_s be the share of defaults occurring in month s . We estimate the average income of all underwater borrowers as:

$$Income^{AllUnderwater} = \sum_s Income_s^{AllUnderwater} w_s. \quad (6)$$

To capture average income of all underwater in months *prior* to default, we compute $Income_t^{AllUnderwater} = \sum_s Income_{s+t}^{AllUnderwater} w_s$ where t is the number of months until default for $t \in \{-12, -11, \dots, 0\}$. Figure 1 shows that this series is essentially flat.

Although our approach relies on comparing the evolution of income prior to default for above and underwater borrowers, we emphasize that our approach does *not* depend on assuming that above water and underwater borrowers have the same degree of financial vulnerability. Table 3 shows that underwater borrowers have slightly higher income levels, bank account balances, and mortgage payment due than above water borrowers. Rather, Assumption 3 instead requires that income declines by the same amount *conditional on a life event*. Table A-2 shows that this does indeed hold for one life event we can reliably observe: unemployment.

4.2 Central Estimate for Strategic Defaults

Our central empirical result—that the evolution of income is virtually indistinguishable for underwater and above water defaulters—is shown in Figure 1. The figure shows that income falls for underwater defaulters nearly as much as for above water defaulters. The gray error bars show the 95 percent confidence interval for the difference between the underwater defaulters and the above water defaulters, centered at the average income drop for the underwater defaulters. This enables visual evaluation of whether the income drop is different in a statistical sense when comparing underwater defaulters and above water defaulters. To provide quantitative estimates, we analyze data for a three month “pre-period” well before default ($t = \{-12, -11, -10\}$) and around the time of default ($t = \{-2, -1, 0\}$). We regress

$$\frac{Income_t}{Payment_{pre}} = \lambda + \kappa \mathbf{1}(LTV > 100) + \underbrace{\gamma \mathbf{1}(t = -2, -1, 0)}_{\text{abovewater drop at default}} + \underbrace{\beta \mathbf{1}(t = -2, -1, 0) \times \mathbf{1}(LTV > 100)}_{\text{difference for underwater}} + \varepsilon \quad (7)$$

where $\overline{Payment}_{pre}$ is the average payment in the pre period, computed separately for above and underwater defaulters. Table 4 column (1) shows that the above water income drop $\hat{\gamma}$ is -0.921. This means that monthly income falls by 92 percent as a share of the monthly payment in the month of mortgage default and the two months prior. Table 4 also shows that $\hat{\beta}$ is 0.051.

Applying the framework from Section 2 to our regression estimates, we find that 94 percent of underwater defaults are causally attributable to life events. Equation (4) requires three inputs to estimate $\alpha_{\text{life event}}$: the change in income for above water defaulters ($\hat{\gamma}$), the change in income for underwater defaulters ($\hat{\gamma} + \hat{\beta}$), and the change in income for all underwater borrowers (which we denote as $\hat{\varphi}$, reported in Table A-3). We estimate that

$$\hat{\alpha}_{\text{life event}} = \frac{(\hat{\gamma} + \hat{\beta}) - \hat{\varphi}}{\hat{\gamma} - \hat{\varphi}} = \frac{-0.921 + 0.051 - (-0.010)}{-0.921 - (-0.010)} = 94\% \quad (8)$$

with a standard error, computed using the delta method, of 1 percent, as shown in Table 5a. The 95 percent confidence interval on the share of underwater defaults causally attributable to negative life events ($\hat{\alpha}_{\text{life event}}$) ranges from 92 percent to 97 percent.

We use the potential outcomes model in Table 2 to interpret this point estimate. Ninety-four percent of underwater defaults have potential outcome types that we categorize as “cash flow” or “double-trigger”, while ($1 - \hat{\alpha}_{\text{life event}} = 1 - 0.94 =$) 6 percent of underwater defaults are “strategic”.

4.3 Robustness for Strategic Defaults

Our finding of very little strategic default is consistent across several specifications and subsamples.

Income Change in the Month of Default Our main quantitative estimate examines the income drop in a three month period around default. Our estimate of the share of strategic default is even smaller (3 percent) if we focus instead on the income drop just in the month of default, as shown in Table 5a.

Distribution of Income Changes Our finding of little strategic default does not hinge on the choice of the mean as a summary statistic. To demonstrate this, the histogram in Figure 2 shows that the entire distribution of the change in monthly income is similar for above and underwater borrowers.²⁰ Quantitative estimates in columns (2)-(4) of Table 4 similarly show that the 25th, 50th, and 75th percentile of the income change distribution are similar for above and underwater borrowers. Finally, event study patterns for different percentiles of the income change distribution are also similar, as shown in Figure A-5.

²⁰Figure A-4 shows the same lesson using the cumulative distribution function.

The distribution of income changes is also useful for evaluating a testable implication of Assumption 3. Income of above water defaulters who experienced a life event could fall *more* than that of underwater defaulters if above water borrowers with smaller income drops used alternative forms of credit to avoid default, while underwater borrowers did not. On the other hand, income of above water defaulters who experienced a life event could fall *less* than the of underwater defaulters if above water borrowers with larger income drops sold their homes quickly (presumably at a heavily-discounted price) before they missed payments. This option is not available to underwater borrowers. Both of these scenarios share a testable implication: the distribution of income drops should exhibit differences in dispersion by home equity, above and beyond any differences in the mean income drop at default. However, the distribution in Figure 2 shows no evidence of such heterogeneity.

One notable feature of Figure 2 is that about one-third of above water defaulters have income *increases*. This may be surprising because Assumption 1 posits that above water defaulters have a life event, which presumably entails a *decrease* in income. This pattern arises because bank account income is a noisy measure of true household income. The increase in income may reflect a severance payment after job loss. It also could reflect the borrower liquidating a retirement account and transferring the funds to her checking account to cover an unexpected expense.²¹

Time Period The finding of little strategic default also does not depend on the time period we study, and in particular holds in 2010 when economists might have expected to see substantial strategic default. Our data are available beginning in January 2007, which means that we have a year of bank account history for defaults which occurred in January 2008 or later. The mortgage delinquency rate peaked in the first quarter of 2010 and Bhutta, Dokko and Shan (2017) show that Google searches for “strategic default” rose massively, peaking in the second quarter of 2010. Figure 3 shows that income declines are similar for above and underwater borrowers in every year from 2008 through 2015. Likewise, Table 5b shows that the estimates of $\alpha_{\text{life event}}$ are similar across different years.

Mortgage Type We find similar patterns across different mortgage types. Figure A-6 plots income before default by home equity separately for borrowers with fixed rate and

²¹The interpretation of our results is the same if some of the income increases reflect liquidation-inducing expense shocks. An alternative scenario is to consider expense shocks that are unobserved, in the sense that they have no effect on bank account income. The interpretation of our results is unchanged if such unobserved shocks have the same relative prevalence for underwater and above water borrowers as observed income shocks. However, if such shocks are more (resp., less) common for underwater borrowers, then our estimates will understate (resp., overstate) the prevalence of strategic default. Finding appropriate proxies for expenditure shocks is a useful direction for future research. The same logic applies to other types of negative life events that may not affect bank account income but can still lead borrowers to miss mortgage payments, such as the onset of a disease like Alzheimer’s.

adjustable rate mortgages, showing similar patterns. Table 5b shows estimates of $\alpha_{\text{life event}}$ for each mortgage type.

Non-recourse State Our results are also not sensitive to whether the default occurs in a non-recourse state. Figure A-7 and Table 5b show that there is no evidence of more strategic default in states with non-recourse mortgage debt, where defaulting may be more financially advantageous. However, this may not be surprising for two reasons. First, Ghent and Kudlyak (2011) reports that deficiency judgments against borrowers in recourse states are exceedingly rare. Second, Guiso, Sapienza and Zingales (2013) reports no difference between borrowers in recourse and non-recourse states in their subjective estimates of the probability that lenders would pursue them after a default.

Definition of Default We observe similar patterns in the data regardless of the definition of default. Instead of dating default based on the number of missed payments by the borrower, it is possible to use the date a lender decides to initiate a foreclosure. Figure A-8 shows that the income patterns before foreclosure are similar for above and below water borrowers. Table 5a shows that our estimate for the share of strategic default using this definition of default (4 percent) is also similar to our baseline estimate.²² Furthermore, although our baseline specification defines default as three missed payments, Figure A-9 shows that the income patterns before default are also similar between above and underwater borrowers when we instead define default as one, two, four, or five missed payments. Table A-5 reports estimates of $\alpha_{\text{life event}}$ for each of these missed payment specifications. In addition, in Appendix F we investigate the path of income after default and find a tight link between income and continued missed payments, regardless of how we define the initial default date.

Heterogeneity by Loan-to-Value Ratio Our results are also not sensitive to the choice of LTV threshold. Figure 4 plots the path of income before default separately for eight different LTV groups. We find large income drops before default across the LTV distribution, from borrowers with substantial positive equity ($\text{LTV} < 60$) to borrowers with substantial negative equity ($\text{LTV} > 180$).

The stability of the income drops across the LTV distribution is surprising relative to prior evidence showing that strategic default is more common for more underwater borrowers. To investigate this further, Table 5b shows estimates of $\alpha_{\text{life event}}$ for LTV bins in increments of 20 percent deep into the LTV distribution. There is little evidence of meaningful strategic default behavior until borrowers have LTVs above 200. Our estimates suggest that the share of strategic default is 11 percent for defaulters with LTVs between 200 and 220 and 15 percent for those with LTVs over 220. However, these groups account for a very small share

²²Table A-4 replicates the regression in equation (7) using foreclosure as the definition of default.

of defaults: only 0.5 percent of defaulters have LTVs above 200.²³

Figure 5 compares our estimates of strategic default ($1 - \hat{\alpha}_{\text{life event}}$) by LTV to those in Bhutta, Dokko and Shan (2017) (henceforth, BDS). Their estimation method uses regional covariates such as unemployment and credit card delinquency to control for negative life events. They find that about 30 percent of all underwater defaults are strategic, and that this share is rising steeply across the LTV distribution. Thus, in quantitative terms, we find significantly less evidence of strategic default than BDS. However, in qualitative terms, our findings are consistent with BDS, who conjecture that improvements in the measurement of life events at the borrower level may lead to smaller estimates of strategic default.²⁴

Measurement Error in Loan-to-Value Ratio Finally, our estimates change little after accounting for measurement error in LTVs using a two-sample instrumental variables approach. Such measurement error arises because the observed home values at the time of default depend on a repeat sales price index. If observed LTV was *completely* noise, then our research design would mechanically (and erroneously) estimate that $\hat{\alpha}_{\text{life event}}$ is 1, regardless of the true value of $\alpha_{\text{life event}}$. The extent to which our estimates are biased depends on the amount of measurement error in observed LTV. Although we do not observe the error in the observed value for any *individual* home, it is straightforward to quantify the *distribution* of errors. This is possible because, using a validation sample from CoreLogic of the subset of homes that actually sell, we can compare observed values to true value at time of sale. We provide details on this data in Section B.3. The availability of such a validation sample is what makes it straightforward to address measurement error in LTVs; in contrast our method in Section 2 must be deliberately agnostic on the distribution of measurement error in life events because no such validation sample exists that could be used to discipline this distribution.

The distribution of errors is approximated well by a Cauchy distribution, as shown in Figure A-11. Let G^* be an unobserved indicator for whether the borrower is truly underwater and LTV be the observed LTV ratio. If we assume that the measurement error in LTV follows a Cauchy distribution, we can construct the probability that an individual borrower is truly underwater ($G^* = 1$) as a function of the borrower’s observed LTV ($P(G^* = 1|LTV)$). For

²³Figure A-10 shows the distribution of LTVs for defaulters in our sample. The small share of borrowers with $LTV > 200$ is not a unique feature of our sample. In the CRISM data, only 1.2 percent of defaulters had $LTV > 200$. This degree of negative equity requires price declines so large that this was rare, even during the Great Recession.

²⁴BDS write that if after controlling for regional covariates that “the remaining unobserved liquidity shocks [life events] correlate with house price shocks” then the contribution of life events “will be even larger than our estimate already suggests.” One way to interpret our estimates jointly is that this omitted variable of life events *is* correlated with house price shocks. Because the JPMCI data have a measure of this omitted variable that can be observed at the borrower level, we learn that strategic default is significantly less prevalent than it appeared to be in prior work.

example, a borrower whose observed LTV is 60 has a 5 percent chance of being actually underwater, and a borrower whose observed LTV is 140 has a 92 percent chance of being actually underwater. In contrast, a borrower with observed LTV of 100 has a roughly 50 percent chance of being underwater.

We estimate that the bias in $\hat{\alpha}_{\text{life event}}$ arising from measurement error in LTV is negligible. If we could observe whether each borrower is truly underwater (G^*), we would estimate

$$\begin{aligned} \frac{\overline{Income}_t}{\overline{Payment}_{pre}} &= \lambda_1 \mathbf{1}(G^* = 0) + \lambda_2 \mathbf{1}(G^* = 1) + \\ &\lambda_3 \mathbf{1}(t = -2, -1, 0) \times \mathbf{1}(G^* = 0) + \lambda_4 \mathbf{1}(t = -2, -1, 0) \times \mathbf{1}(G^* = 1) + \varepsilon. \end{aligned} \quad (9)$$

Because we do not observe G^* , we use an adjustment procedure akin to two-sample instrumental variables, where the first stage sample is from CoreLogic and the second stage sample is from JPMCI. Using $P(G^* = 1|LTV)$ as constructed in the previous paragraph, we can feasibly estimate

$$\begin{aligned} \frac{\overline{Income}_t}{\overline{Payment}_{pre}} &= \lambda_1 P(G^* = 0|LTV) + \lambda_2 P(G^* = 1|LTV) + \\ &\lambda_3 \mathbf{1}(t = -2, -1, 0) \times P(G^* = 0|LTV) + \lambda_4 \mathbf{1}(t = -2, -1, 0) \times P(G^* = 1|LTV) + \varepsilon. \end{aligned} \quad (10)$$

We calculate that the income drop before default for the truly above water is 91.6 percent (compared to 92.1 percent for those observed above water) and for the truly underwater is 87.9 percent (compared to 87.0 percent for those observed underwater).²⁵ Then, to convert the regression estimates to an estimate of the role of life events, we again use equation (4): $\hat{\alpha}_{\text{life event}} = \frac{\hat{\lambda}_4 - \hat{\varphi}}{\hat{\lambda}_3 - \hat{\varphi}}$. Table 5a shows that $\hat{\alpha}_{\text{life event}}$ rises from 0.94 in our main specification to 0.96 when adjusting for this issue.²⁶ The estimates are similar even when we consider more extreme adjustments which are likely to overstate the extent of measurement error in LTV, such as assuming that a measure of the LTV error in the Chase data is twice as large as the error in the CoreLogic validation data or assuming that all homes in default sell at a 9 percent discount (see Table A-6).²⁷

²⁵We report measurement error-corrected estimates analogous to equation (7) in Column 5 of Table 4.

²⁶Thus, there is a slight increase in $\hat{\alpha}$. One might have expected that $\hat{\alpha}$ would fall after adjusting for measurement error. However, the measurement error correction procedure upweights borrowers with LTV < 60 and borrowers with LTV > 140 relative to borrowers with LTV closer to 100. Because income drops for borrowers with LTV < 60 and LTV > 140 are slightly more similar than the income drops for all underwater borrowers versus all above water borrowers, the estimate of $\hat{\alpha}$ is slightly larger.

²⁷Nine percent is an upper bound on the plausible discount for defaults in the Chase data. Giacoletti (Forthcoming) estimates that *foreclosed* homes sell at an average discount of 5-9 percent relative to their

The intuition for why the estimates change so little after accounting for measurement error can be understood by revisiting the income drop by LTV gradient shown in Figure 4. Two aspects of the figure are worth noting. First, if there was substantial strategic default and LTV was measured with error, then we would expect to see income drops that are progressively smaller in magnitude for higher LTV bins. Instead, the average income drop is very similar, regardless of LTV. Second, this pattern holds even where the borrower has an LTV of less than 60, such that they are almost certainly above water, or an LTV of greater than 140, such that they are almost certainly underwater. Because the income drops are so similar across measured LTV groups, there is little change in our estimates of $\hat{\alpha}_{\text{life event}}$ from accounting for measurement error in LTV.

4.4 Who are the Strategic Defaulters?

The prior section shows evidence of very little or no strategic default across a wide variety of specifications. A natural question is whether our method can detect evidence of more substantial strategic behavior in sub-populations where we might expect this behavior to be more prominent. In fact, our method does detect more evidence of strategic behavior in one sub-population: those that miss three straight mortgage payments. However, the aggregate magnitude of strategic default implied by this subgroup analysis is small.

The analysis of borrowers who miss three straight mortgage payments is inspired by Mayer et al. (2014), which notes that a borrower who decides to strategically default will stop making payments once and for all.²⁸ Therefore, if there is evidence of strategic default, it should manifest within this sub-population. At the same time, missing three straight mortgage payments may simply reflect a borrower facing a severe economic shock.

We do indeed see stronger evidence of strategic default in this sub-group. Figure A-12 shows larger income declines for above water defaulters than for underwater defaulters in the subsample that misses three consecutive mortgage payments. In Table 5b, we estimate that 12 percent of underwater defaults are strategic for this subsample.

Another group that we might expect to disproportionately default strategically is investors (Albanesi, De Giorgi and Nosal 2017). Although our sample is representative relative

index-based estimate, but many defaulters are not foreclosed on.

²⁸Other papers that study consecutive missed payments include Keys et al. (2012), Bradley, Cutts and Liu (2015) Experian and Wyman (2009) and Tirupattur, Chang and Egan (2010). Keys et al. (2012) measure the share of mortgage defaults that transition straight from 60 days past due to 180 days past due in four months, while remaining otherwise current on all non-HELOC revolving debt. We extend the analysis in Keys et al. (2012) forward through 2015 using the CRISM data and show the results in Figure A-13. We find that 16.3 percent of underwater borrowers meet their definition of sequential missed payments, while 10 percent of abovewater borrowers meet this definition. The excess sequential default rate for underwater borrowers is 6.3 percent. If we interpret this as an alternative estimate of the prevalence of strategic default, it falls within the confidence interval of our central estimate.

to external benchmarks in terms of the share of *self-declared* investors (see Table A-1), we cannot directly speak to the types of investors found to be more strategic in prior work. In particular, Elul, Payne and Tilson (2021) and Piskorski, Seru and Witkin (2015) document that self-declared investors behave similarly to other borrowers. Within the investor group, it is *fraudulent* investors who appear more strategic. These are borrowers who claim to be owner-occupants but who in fact have multiple first liens. We do not observe first liens outside of our data set and so cannot identify such borrowers. If such fraudulent investors are disproportionately missing from our data, then our estimates may understate the population-wide prevalence of strategic default. However, Elul, Payne and Tilson (2021) document that such investors were broadly distributed across GSE, portfolio, and privately securitized loans. Moreover, because this sub-population is small, we note that our quantitative estimates of the overall share of strategic default are consistent with meaningful strategic default among these borrowers. Elul, Payne and Tilson (2021) reports that fraudulent investors account for only 10 percent of all defaults.

4.5 Causal Impact of Negative Equity

4.5.1 Research design

We now estimate the causal impact of negative equity on default in order to identify $\alpha_{\text{negative equity}}$ and complete the decomposition between the three theories of default. Palmer (2015) estimates this causal impact for subprime borrowers using an instrument based on long-run city-level variation in house price cyclicalities. Guren et al. (2021) further develops this instrument and uses it to estimate housing wealth elasticities. Below, we apply the cyclicalities method to estimate the causal impact of negative equity. We find that the estimates for subprime borrowers in Palmer (2015) extend to a broader sample of borrowers using Chase and CRISM data.

We estimate the impact of negative equity on default in a Cox proportional hazard model:

$$h_{i,t,m,j,r} = h_0(t) \exp \{ \psi_j + \xi_{r,m} + \delta \mathbf{1}(LTV_{i,t} > 100) + X_i' \theta \}, \quad (11)$$

where $h_{i,t,m,j,r}$ denotes the default hazard rate for borrower i , at loan age t months since origination, in calendar month m (i.e. “February 2010”), CBSA j , and Census region r . $h_0(t)$ denotes the nonparametric baseline default hazard, ψ_j denotes CBSA fixed effects, $\xi_{r,m}$ denotes region-by-calendar-month fixed effects, $\mathbf{1}(LTV_{i,t} > 100)$ is an indicator equal to 1 if borrower i has negative equity in month t , and X_i' is a vector of borrower and loan characteristics measured at origination.²⁹

²⁹Borrower and loan characteristics include an indicator variable for whether the loan is a balloon mortgage,

The coefficient of interest in equation (11) is δ , which measures the proportional increase in the default hazard associated with being underwater. There are two main challenges with identifying δ . The first challenge is the potential endogeneity of negative equity. This can arise from reverse causality (defaults cause price declines that push borrowers underwater) or omitted variables (negative life event shocks like local unemployment shocks can directly decrease house prices by reducing housing demand and also directly increase default rates). The second challenge is measurement error in negative equity. Both challenges can be addressed with a two-stage instrumental variable approach.

We use the cyclical instrument from Palmer (2015) and Guren et al. (2021). The instrument is an interaction between a measure of each city’s average sensitivity to price cycles over a long time period and either calendar-month indicator variables (as in Palmer 2015) or regional price changes (as in Guren et al. 2021). The key idea behind the instrument is that house prices in some cities are systematically more sensitive to regional price cycles than are house prices in other cities. Palmer (2015) shows that city-level cyclical instrument is unrelated to housing market shocks unique to the 2000s price cycle (such as the subprime credit expansion), and Guren et al. (2021) extends this instrument to use the component of variation that is orthogonal to many potential confounds.

The city-level cyclical instrument is highly predictive of LTVs, indicating a strong first stage. We regress LTV on the same fixed effects and controls as in (11) as well as the interaction between the sensitivity measure and house price changes, i.e.

$$LTV_{i,t,m,j,r} = \Psi_j + \Xi_{r,m} + \rho\Gamma_j\Delta P_m + X_i'\Theta + \eta_{i,t,m,j,r} \quad (12)$$

where Γ_j is the city level house price sensitivity measure reported in Guren et al. (2021), ΔP_m is the log annual change in regional house prices, and $\eta_{i,t,m,j,r}$ is the error term.³⁰ The first stage is presented visually in Figure A-14 using a residualized binscatter. The F-statistic is high (over 100). Intuitively, borrowers who live in cities more sensitive to price cycles are more likely to have higher LTVs when aggregate house prices have fallen.

We use the control function approach in order to instrument in the non-linear hazard model setting of equation (11), following Palmer (2015) and Imbens and Wooldridge (2007). This involves adding the estimated residuals $\hat{\eta}_{i,t,m,j,r}$ from (12) as controls when estimating

the interest rate at origination, an indicator variable for whether the loan was interest-only at origination, the origination LTV, an indicator for whether the home is the borrower’s primary residence, an indicator for whether the borrower is subprime, and an indicator for whether the loan is an ARM.

³⁰Following Guren et al. (2021) we use the house price index from Freddie Mac at the census region level, and convert it to a real house price index using the GDP deflator. We also consider an alternative first-stage specification which interacts Γ_j with calendar-month fixed effects rather than regional house price changes, which we discuss below.

equation (11).

4.5.2 Causal impact of negative equity on default: estimate and robustness

Table 6 shows our results. Panel (a) shows estimates of the default hazard function in equation (11) for borrowers in the Chase sample. Column (1) shows the most parsimonious specification, without individual controls and without instrumenting for negative equity. The coefficient on the underwater dummy shows an estimate of $\hat{\delta} = 1.51$. Controlling for borrower and loan characteristics in column (2) reduces this estimate slightly to 1.31.

Our preferred specification is in column (3), which instruments for negative equity by including the LTV fitted residuals ($\hat{\eta}$). We find that $\hat{\delta} = 0.36$, which means that negative equity causally increases the default hazard by 43 percent ($\exp(\hat{\delta})$). The statistical significance of the LTV residuals suggests that the naive underwater indicator is indeed endogenous. Column (4) shows that we find similar effects when we include origination-year fixed effects and specify the instrument as the interaction between city-level sensitivity and calendar-month fixed effects, as in Palmer (2015). This specification leads to an estimate of $\hat{\delta} = 0.32$.

Our estimates are similar when we use a different data source and are also similar to prior estimates for specific sub-populations. Panel (b) replicates the specifications from panel (a) in the CRISM dataset, which capture mortgages from a broad range of lenders. We find IV estimates of $\hat{\delta}$ between 0.24 and 0.27 in columns 3 and 4. The estimates for borrowers from both panels in Table 6 (which capture both prime and subprime borrowers) are in line with the estimates for subprime borrowers in Palmer (2015), which finds an effect of negative equity equivalent to a $\hat{\delta}$ of 0.28.

4.5.3 Interpretation: share of defaults caused by negative equity

Our preferred estimate is that the share of underwater defaults caused by negative equity ($\alpha_{\text{negative equity}}$) is 30 percent. Combining the causal impact of negative equity with the formula from equation (5), $\alpha_{\text{negative equity}} = 1 - E \left[\frac{Y(T^*,1)}{Y(T^*,0)} | G = 1 \right]^{-1} = 1 - \exp(-\delta)$. At our preferred estimate of $\hat{\delta} = 0.36$, $\hat{\alpha}_{\text{negative equity}} = 30$ percent. Estimates of $\hat{\alpha}_{\text{negative equity}}$ from each specification are reported in the bottom rows of Table 6, along with standard errors calculated using the delta method. The IV estimates are similar across two different ways of constructing the instrument and two different datasets, ranging from 22 percent to 30 percent.

We would reach a similar conclusion if instead of directly estimating the causal impact of negative equity on default, we combined our framework with an alternative research design in prior work. Gupta and Hansman (2021) estimates the causal impact of negative equity on default using a natural experiment arising from the unexpected divergence between indexes that determine interest rates for option ARM mortgages. This natural experiment generates

exogenous variation in mortgage balances between otherwise similar borrowers. They find that negative equity causally increases default probabilities by about 30 percentage points and that the underwater default probability is about 60 percent. Plugging these numbers into equation (5) delivers an estimate of $\hat{\alpha}_{\text{negative equity}} = 50$ percent for their sample of option ARM borrowers.

4.6 Decomposition: Strategic vs Cash-Flow vs Double-Trigger Defaults

Using $\hat{\alpha}_{\text{life event}}$ from Section 4.2 and $\hat{\alpha}_{\text{negative equity}}$ from Section 4.5 allows for a full partition between the three theories of default. Our central estimate of $\hat{\alpha}_{\text{negative equity}}$ is that eliminating negative equity would eliminate 30 percent of underwater defaults. Interpreted through the lens of the potential outcomes model in Table 2, this suggests that 30 percent of underwater defaults have the potential outcome types that we categorize as “strategic” or “double-trigger,” while the remaining ($1 - \hat{\alpha}_{\text{negative equity}} = 1 - 0.3 =$) 70 percent of underwater defaults are “cash-flow” defaults. Combined with our prior estimate from $\hat{\alpha}_{\text{life event}}$ that 6 percent of underwater defaults are strategic, this means 24 percent are “double-trigger”.

The finding that double-trigger behavior accounts for about one-quarter of underwater defaults suggests an important role played by the *interaction* between negative life events and negative equity in driving default. Indeed, there is strong prior evidence to support double-trigger behavior. For example, seminal work by Gerardi et al. (2018) provides the first direct empirical evidence for this type of default. Their study shows that, conditional on an observable life event, the probability of default is higher for borrowers with negative equity.³¹ Double-trigger defaults can arise mechanically because it is more difficult for underwater borrowers to avoid default by selling or borrowing against their homes, or behaviorally if underwater borrowers are less likely to prioritize their mortgage payments after a life event (Chan et al. 2016).

However, the findings also show that pure cash-flow defaults, driven *entirely* by negative life events, are also important. This stands in sharp contrast to the standard view in the literature. Our preferred estimate is that the commonly-neglected cash-flow defaults account for 70 percent of all underwater defaults. Even if one were to take the extreme assumption that the entire cross-sectional relationship between negative equity and default is causal, we still find that a significant fraction of defaults are driven entirely by cash flow. This specification, from column (1) of Table 6, implies that the pure cash-flow channel accounts for about one-quarter of all underwater defaults.

³¹See also Cunningham, Gerardi and Shen (2020) for more evidence of double-trigger behavior.

5 Re-analysis of Survey Data

Our empirical findings may be surprising because some prior work estimates that between 30 and 70 percent of Great Recession defaults are strategic (Gerardi et al. 2018; Guiso, Sapienza and Zingales 2013; Bhutta, Dokko and Shan 2017). It is natural to wonder whether our much lower estimate arises from differences in data, differences in the definition of strategic default, or differences in estimation methodology. By applying our methodology of using above water defaulters as a comparison group to survey data on income and default in the PSID, we can distinguish between these hypotheses.

The PSID has two additional benefits relative to the JPMCI data. First, it captures a sample from all bank accounts and mortgage servicers (rather than from the universe of people who have checking accounts at, and a mortgage serviced by, one bank). In addition to providing a comprehensive view of a household’s financial circumstances, this also helps address a concern about whether underwater borrowers hide some of their income from the bank that services their mortgage.³² Second, it captures borrowers’ *perceived* LTV, which is the decision-relevant measure of LTV from the borrower’s perspective, providing an alternative way to circumvent possible concerns about measurement error in LTV described above. However, the JPMCI data have monthly income (rather than once every two years) and there are 500 times as many defaults in the JPMCI data as in the PSID (which has only 244 households that default on their mortgages).

We anchor our analysis on a definition of strategic default from pioneering work by Gerardi et al. (2018, henceforth GHOW) and Guiso, Sapienza and Zingales (2013). Guiso, Sapienza and Zingales (2013) surveys a representative sample of US households, about one-third of whom report that they know someone who has defaulted on their mortgage. They ask this subsample whether their defaulting acquaintances “could afford to pay the monthly mortgage.” GHOW take a similar approach, though instead of asking acquaintances, they use self-reported information from the PSID. They classify a mortgage as affordable when a borrower “*can pay* their mortgage without reducing consumption from its predefault levels.” We focus our analysis on the PSID because it includes information on whether borrowers are above or below water, and so we can implement our comparison group approach.

To measure mortgage affordability empirically in the PSID, GHOW examines the dis-

³²This concern might arise because of two different types of borrower misperception. First, if a borrower incorrectly believes that assets at the mortgage-servicing bank are more likely to be seized than assets at another bank. Second, manipulation could also arise if the borrower incorrectly believes that lower bank account income will lead to a more generous mortgage modification. However, such beliefs are inaccurate. In fact, seizure can only occur after a foreclosure is completed and a deficiency judgment has been rendered. In addition, the bank’s publicly-available mortgage assistance form asks for paystubs and income as reported to tax authorities; it does not ask for bank account income.

tribution of disposable income (income y minus non-housing consumption c) net of housing expenses m . The idea behind this analysis is that if an underwater borrower’s mortgage is affordable (i.e., “available resources” = $y - c - m > 0$) and yet they default, then they must be defaulting strategically.

We re-analyze the prevalence of underwater strategic default in PSID, adding above water defaulters and all underwater borrowers as comparison groups. Our PSID analysis is in the spirit of the previous sections of this paper, but directly applying the method from Section 2 yields statistically imprecise results.³³ Instead, we use the *can pay* definition of strategic default to enable comparability to the prior literature.

Figure 6 plots the cumulative distribution function of available resources for above water defaulters, underwater defaulters, and borrowers who are not in default. Consistent with a role for negative life events, available resources are much lower for defaulters than non-defaulters. In contrast, when we look *within* defaulters to compare above and underwater borrowers, the distribution of available resources appears to be the same. Thus, the results in the PSID echo the conclusions of Figure 2 (and Figure A-4) in finding very little evidence of strategic default.

The key reason why we find little strategic default when prior work found substantial strategic default is the adjustment for measurement error, rather than differences in data source or differences in the definition of strategic default. For example, Figure A-15a shows that 39 percent of underwater defaulters meet the *can pay* definition and one might conclude that these 39 percent of underwater defaults are strategic. Yet Figure A-15a shows that 37 percent of above water defaulters also meet the *can pay* definition in a sample that has no motive to default strategically. The difference between the share of defaults labeled *can pay* for above water versus underwater is not statistically significant (p-value of 0.82). We conclude that income and consumption obligations are difficult to measure at the household level.³⁴ Not accounting for measurement error may lead researchers to understate the importance of life events, as noted by Gyourko and Tracy (2014).

This null finding appears to be a fundamental feature of the joint distribution of available resources and home equity among defaulters rather than a result of one specific definition

³³Our method is underpowered in the PSID. We are unable to reject the hypothesis that *no* defaults are strategic and also that *all* defaults are strategic. This is not surprising. In general a noisy variable on the left-hand side of a regression creates larger standard errors, but not bias. Our method, which places a noisy measure of treatment on the left-hand side, sacrifices precision to avoid bias. This is not a problem for a large administrative dataset, but it is a problem for the PSID, where there are only 244 mortgage defaulters.

³⁴For example, precise measurement of ability to pay at the time of mortgage default may be difficult if the borrower experienced an income shortfall in one month (even though calendar year income was sufficient to cover the annual mortgage payment). It also might be difficult if the borrower underreported consumption, as people are known to do in recall surveys (Passero, Garner and McCully 2014)

of mortgage affordability. GHOW also analyzes an alternative definition of mortgage affordability, which examines how many people would “need to reduce consumption below subsistence levels to remain current on their mortgage.” Using this measure of affordability, Figure A-15a shows that the share of defaults labeled as strategic is quantitatively similar for above and underwater borrowers (p-value of 0.23). Further, Figure A-16a shows that the entire distribution of $y - c_{subsistence} - m$ is similar. In addition, Figure A-15a shows that the share of defaults with an income increase is similar for above and underwater borrowers in the PSID and in the JPMCI data. Because above and underwater borrowers have the same distribution of available resources and the same change in income around default, we conclude that there is little evidence of strategic default in the PSID.³⁵

Although the distributions of available resources among defaulters in the PSID shows little evidence of strategic behavior, the patterns do show clear evidence of double-trigger behavior. As documented in GHOW, conditional on any of the measures of affordability discussed above, underwater borrowers are more likely to default than above water borrowers. This suggests that negative equity *interacted* with affordability is indeed an important driver of default.

6 Implications for Modeling Mortgage Default

In this section, we show that our empirical results contrast with predictions from standard structural models of mortgage default. However, we show that an extension where default has a high utility cost can bring the standard model in line with our new empirical moments. This reconciliation may provide a blueprint for a wide class of macro finance models where borrower default decisions play a central role.

To provide a model-based comparison to our empirical moments we use the modern benchmark quantitative model of mortgage default first developed in Campbell and Cocco (2015, henceforth CC). This model is ideal for assessing whether existing structural models can match our empirical moments because it is the first to integrate strategic motives based on option-value theory and cash-flow motives based on realistic income risk. An extensive prior literature uses option-value theory to model default decisions based on house price risk (Epperson et al. 1985; Foster and Van Order 1984; Deng, Quigley and Van Order 2000). CC adds idiosyncratic income risk to the option-value model in order to quantify the relative contributions of negative equity and cash-flow motives to the default decision. Relative to

³⁵These figures define underwater as LTV greater than 100, to be consistent with the rest of our paper. Figures A-15b and A-16b present the same analysis using an LTV cutoff of 90, which is the cutoff used in GHOW.

the prior real-option models, it is exactly this new type of model incorporating cash-flow motives that has the best chance of matching our empirical results.

A secondary benefit of comparing our empirical results to predictions from a structural model is that the model provides a framework for relaxing one of the common assumptions in empirical work. Specifically, in Section 2 we assumed that each of the treatments was binary (e.g., that each household has a negative life event, negative home equity, or both). Put otherwise, the potential outcome function $Y(T^*, G)$ uses binary T^* and binary G . Similarly, the prior empirical work discussed in Section 5 seeks to classify borrowers in a binary fashion as either experiencing a life event or not. In contrast, structural models allow for a cash-flow shock with many possible realizations, such that it is possible to generate an agent’s policy function $Y(T^*, G)$ with discrete T^* .³⁶

The CC model has a novel prediction that cash-flow motives dominate for households with slightly negative equity while strategic motives dominate for households with substantially negative equity. Rather than classifying borrowers as strategic in a binary fashion as in prior empirical work, they instead report summary statistics by LTV bin in Figure 2 of their paper.³⁷ The figure shows that among households with LTV between 100 and 120, the income of defaulters is substantially lower than the income of non-defaulters. CC explains that this pattern emerges because short-term cash flow considerations drive the default decisions of moderately underwater borrowers:

As house prices decline, households with tightly binding borrowing constraints will default sooner than unconstrained households, because they value the immediate budget relief from default more highly relative to the longer-term costs.

In contrast, for households with LTV above 150, the income of defaulters is the same as non-defaulters. For these households, immediate budget constraints are less important. Instead, negative equity drives the default decision. In this way, the CC model captures the intuition of how both negative life events that require immediate budget relief and strategic motives contribute to the default decision.

This novel prediction from the CC model is also apparent when we replicate our empirical specification within the model. This requires two extensions. First, instead of comparing the level of income of defaulters to that of non-defaulters, our approach relies on calculating the within-borrower change in income prior to default. We show how this statistic varies

³⁶The shock is not continuous in the CC model because the simulation method uses Gaussian quadrature.

³⁷This figure evaluates the joint distribution of income, home equity and default for mortgagors with adjustable-rate mortgages. It shows four scenarios with varying levels of income risk and initial yield rates. CC write that the lessons from this figure are “most visible in Panel D”, which is the scenario with high income risk and high initial yield. We use this scenario throughout our analysis. Figure A-17 shows that we can replicate the summary statistics from CC’s Table II panel D.

with LTV both in the baseline CC model and in our data in Figure 7a.³⁸ Second, although above water borrowers do not default on their mortgages in the CC model, they do sell their houses and then terminate their mortgages by prepaying them. Thus, to provide a model counterpart to the empirical income drop for above water defaulters, we compute the income drop for above water prepayers in the model.

Figure 7a shows a substantial income drop both at prepayment for above water borrowers and at default for borrowers with LTV between 100 and 120. The intuition for why an income drop precedes both types of mortgage termination in the model is that they both cause borrowers to lose out on an investment that requires upfront liquidity but has long-term positive expected returns. Thus these borrowers are only likely to terminate if they have suffered a liquidity shock. For expositional simplicity, we use the term “default” to describe both types of mortgage termination. However, the central conclusions from this comparison of model to data are unchanged if we only focus on the income drop gradient for the underwater defaulters and ignore the group of above water prepayers.

Two main conclusions emerge from comparing the model’s predictions to the data. First, for defaulters with LTV less than 120, the baseline model’s predictions are remarkably similar to our empirical results. Figure 7a shows that these borrowers suffer a substantial income drop before default in both the model and the data. The similarity between model and data is particularly surprising because individual-level estimates of income losses prior to default were not available as targets for the development of the model. In addition, the model predicts that borrowers in this range exhaust their liquid assets before defaulting. This accords with empirical evidence that borrowers who default have virtually no liquid assets (see Table 3, which shows mean checking account balances equal to less than two weeks of income, and also Gerardi et al. 2015). This similarity of the model’s predictions to the data indicates that the model successfully captures defaults triggered by negative income shocks, which themselves may arise because of a life event.

Second, in contrast to the tight match between model and data for moderately underwater borrowers, we find a sharp divergence for borrowers with substantial negative equity. For borrowers with LTV ratios above 120, the model’s predictions converge to the standard option-value framework, where defaults are driven by negative equity rather than by individual cash flow. Figure 7a shows that the model predicts a drop in income (as a share of mortgage payment due) of 15 percent or less, while in the data the drop is five times as large. A similar divergence holds when considering liquid assets. In the model, we find that these

³⁸The time interval in CC’s model is annual and it takes a few years for a substantial negative permanent income shock to accumulate because CC use Gaussian quadrature (as is conventional in this literature). We therefore calculate the income change in the model over a four-year time horizon to allow for meaningful negative income shock realizations relative to baseline.

defaulting borrowers have mean liquid assets equal to a bit over one *year* of income. This contrasts with the finding in Table 3 that underwater checking account balances prior to default are equal to less than two *weeks* of income, just like the above water borrowers. Table A-8 further shows that, at the time of default, the entire distribution of checking account balances for underwater borrowers is similar to that for abovewater borrowers and suggests that they have very little liquidity available to cover a mortgage payment. Intuitively, as LTV increases in the model, more borrowers prefer to default because their homes are a bad financial investment (as in Foster and Van Order 1984), regardless of whether they need what CC calls “immediate budget relief.” In contrast, the decisions of borrowers in the data appear less influenced by the value of the house as a financial asset than is expected in the model. Instead, a substantial income drop and exhaustion of assets precedes default even for deeply underwater borrowers.

Despite this divergence between model and data for deeply underwater borrowers, an extension to the baseline model offers a potential reconciliation. In the baseline model, deeply underwater borrowers default in the absence of immediate cash-flow motives because the long-term financial benefits are large while the costs are small. But in practice, the costs of defaulting may also be large. For example, defaulting may impose financial costs through reduced access to credit, or non-financial costs due to borrowers’ attachment to their current home, a fear of social stigma, or a moral aversion to default. The CC model builds in the possibility that these costs are important through an extension that allows for a utility cost of default, which CC calls *Stigma*. However, the paper explains that “the main difficulty with this extension of our model is determining an appropriate value of *Stigma*.” Indeed, prior work has estimated a wide range of default costs, from as low as a 1.5 percent decrease in the constant equivalent-consumption stream to as high as a 70 percent decrease (Kaplan, Mitman and Violante 2020; Hembre 2018; Schelkle 2018; Laufer 2018).

We propose to use the income drop before default as a new moment to discipline this parameter. Intuitively, the reluctance of borrowers to default on a substantially underwater asset in the absence of immediate budgetary pressure is informative about how costly they perceive this default to be. We therefore estimate the utility cost that minimizes the distance between the model’s predicted income declines for underwater defaulters and the income declines we actually observe in the data. The best fit is that defaulting imposes a one-time utility loss equal to a 25 percent decrease in the constant-equivalent consumption stream or \$100,000 in present value terms (see Appendix G for details). This is a very high cost. We note that this utility cost includes behavioral and moral factors; for example, Guiso, Sapienza and Zingales (2013) reports that 82 percent of survey respondents believe that strategic default is morally wrong. The high cost is also consistent with emerging evidence

that foreclosures substantially damage family outcomes (Diamond, Guren and Tan 2020).

Under this alternative parameterization, the model is able to closely match our new empirical moments. This is shown visually in Figure 7b. When defaulting is costly, borrowers in the model only exercise their default option when they need short-term budgetary relief. This generates income drops before default consistent with our empirical results. Thus, a high cost of default provides a plausible micro-foundation for the behavior we observe empirically within a benchmark optimizing framework. Furthermore, recent evidence suggests that households are responsive to variation in default costs. O’Malley (2021) finds that default rates rose when default costs were cut dramatically by a legal ruling in Ireland that completely eliminated repossession risk for a subset of borrowers.

This approach to reconciling model and data may provide a blueprint for a wide class of macro finance models where borrower default decisions play a central role. For example, models with endogenous borrower default have been used recently to inform questions about macroprudential regulation, the origins of the 2008 financial crisis, bankruptcy and foreclosure policy, and optimal mortgage security design.³⁹ These types of models need to take a stand on what triggers borrower default. Our empirical results suggest that it is crucial to build in mechanisms that lead life events such as cash-flow shocks to be a necessary condition for default. Our parameterization of the CC model demonstrates that one specific way to achieve this is to incorporate a large utility cost of defaulting. More generally, regardless of exactly how the default decision is modeled, models with endogenous borrower default might seek to target large income drops before default and low assets at the time of default. In Tables A-7 and A-8, we provide moments on the joint distribution of income, assets, and home equity, which may be useful as a target for such models. Incorporating realistic default behavior triggered by negative events into macroeconomic models is an exciting topic for future work.

7 Conclusion

This paper asks a simple question: why do borrowers default on mortgages? The literature has focused on two candidate triggers for default: negative equity and negative life events. However, despite longstanding interest, it has remained difficult to distinguish between these triggers in part because it is difficult to precisely measure life events. We propose a new method to address this measurement error problem using a comparison group of borrowers whose default must have been caused by a negative life event: borrowers with

³⁹For example, see Corbae and Quintin (2015), Mitman (2016), Kaplan, Mitman and Violante (2020), Guren, Krishnamurthy and McQuade (2019), Campbell, Clara and Cocco (Forthcoming), Greenwald, Landvoigt and Van Nieuwerburgh (2021), Diamond and Landvoigt (2019), and Garriga and Hedlund (2020).

positive home equity. For these borrowers, negative equity cannot be the cause of their default. We implement this method in a new high-frequency dataset linking income and mortgage default.

Our central finding is that only 6 percent of underwater defaults are caused exclusively by negative equity, much less than previously thought. Furthermore, we find that 70 percent of underwater defaults are driven solely by negative life events. Moreover, because even at the peak of the housing crisis at least 40 percent of defaults were by above water borrowers (whose defaults must all be driven by negative life events), the fraction of *all* defaults accounted for by negative life events is even greater. Although our finding contrasts sharply with predictions from standard models, we show that it can be rationalized in models with a high private cost of mortgage default. This reconciliation between model and data may provide a blueprint for general macro finance models where borrower default decisions play a central role.

An additional contribution of this paper is that it may be helpful in designing policies to address mortgage default. If default is driven exclusively by negative equity, then forgiving mortgage principal may be necessary to prevent default. However, principal forgiveness is expensive for lenders and, if the policy is subsidized, for taxpayers. On the other hand, if negative life events are instead a necessary condition for most defaults and if reducing payments temporarily (e.g. for up to a few years) would eliminate many of the defaults caused by negative life events, then temporary payment reduction is a dominant strategy for lenders and policymakers seeking to reduce defaults. Compared to principal forgiveness, payment reduction can be accomplished much more cheaply, for example by offering forbearance or mortgage term extensions, which provide immediately liquidity while leaving the principal balance unchanged (Campbell, Clara and Cocco Forthcoming).

This logic has support in the prior literature, which has shown using within-contract variation among borrowers who received a mortgage modification that reducing payments is more effective than forgiving principal at preventing default (Scharlemann and Shore 2019; Ganong and Noel 2020). However, by construction, the prior literature does not study borrowers excluded from mortgage modifications—which often have stringent eligibility criteria designed to exclude strategic defaulters—and borrowers who are deeply underwater. In contrast, the method presented in this paper allows for estimates of strategic default and cash-flow induced default for *all* borrowers, suggesting that the policy prescription in favor of reducing payments extends to a much broader set of borrowers than previously understood.

This prescription may be useful in future crises. The number of mortgages with missed payments soared during the Covid-19 recession—which was accompanied by an obvious, widespread negative life event—and nearly all lenders have offered immediate liquidity via

forbearance (Cherry et al., 2021). If house prices fall again while unemployment remains high, policymakers may again wonder whether they should tackle negative equity through expensive principal reductions. Our results suggest that focusing on liquidity may be sufficient for nearly all borrowers.

References

- Albanesi, Stefania, Giacomo De Giorgi, and Jaromir Nosal.** 2017. “Credit Growth and the Financial Crisis: A New Narrative.” National Bureau of Economic Research Working Paper 23740.
- Bayer, Patrick, Fernando Ferreira, and Stephen L. Ross.** 2016. “The Vulnerability of Minority Homeowners in the Housing Boom and Bust.” *American Economic Journal: Economic Policy*, 8(1): 1–27.
- Beraja, Martin, Andreas Fuster, Erik Hurst, and Joseph Vavra.** 2019. “Regional Heterogeneity and the Refinancing Channel of Monetary Policy*.” *The Quarterly Journal of Economics*, 134(1): 109–183.
- Bernanke, Ben S.** 2008. “Mortgage Delinquencies and Foreclosures.”
- Bernstein, Asaf.** 2019. “Negative Home Equity and Household Labor Supply.” Social Science Research Network SSRN Scholarly Paper ID 2700781, Rochester, NY.
- Bhutta, Neil, Jane Dokko, and Hui Shan.** 2017. “Consumer Ruthlessness and Mortgage Default during the 2007 to 2009 Housing Bust.” *The Journal of Finance*, 72(6): 2433–2466.
- Black, Sandra E., Paul J. Devereux, Petter Lundborg, and Kaveh Majlesi.** 2019. “Poor Little Rich Kids? The Role of Nature versus Nurture in Wealth and Other Economic Outcomes and Behaviours.” *The Review of Economic Studies*, 87(4): 1683–1725.
- Boar, Corina, Denis Gorea, and Virgiliu Midrigan.** 2020. “Liquidity Constraints in the U.S. Housing Market.” National Bureau of Economic Research Working Paper 23345.
- Bradley, Michael G., Amy Crews Cutts, and Wei Liu.** 2015. “Strategic Mortgage Default: The Effect of Neighborhood Factors.” *Real Estate Economics*, 43(2): 271–299.
- Campbell, John Y., and João F. Cocco.** 2015. “A Model of Mortgage Default.” *The Journal of Finance*, 70(4): 1495–1554.
- Campbell, John Y., Nuno Clara, and João F. Cocco.** Forthcoming. “Structuring Mortgages for Macroeconomic Stability.” *The Journal of Finance*.
- Campbell, Tim S., and J. Kimball Dietrich.** 1983. “The Determinants of Default on Insured Conventional Residential Mortgage Loans.” *The Journal of Finance*, 38(5): 1569–1581.
- Chan, Sewin, Andrew Haughwout, Andrew Hayashi, and Wilbert Van Der Klaauw.** 2016. “Determinants of Mortgage Default and Consumer Credit Use: The Effects of Foreclosure Laws and Foreclosure Delays.” *Journal of Money, Credit and Banking*, 48(2-3): 393–413.
- Cherry, Susan, Erica Xuewei Jiang, Gregor Matvos, Tomasz Piskorski, and Amit Seru.** 2021. “Government and Private Household Debt Relief during COVID-19.” National Bureau of Economic Research Working Paper 28357, Cambridge, MA.
- Chodorow-Reich, Gabriel.** 2014. “The Employment Effects of Credit Market Disruptions: Firm-level Evidence from the 2008–9 Financial Crisis.” *The Quarterly Journal of Economics*, 129(1): 1–59.

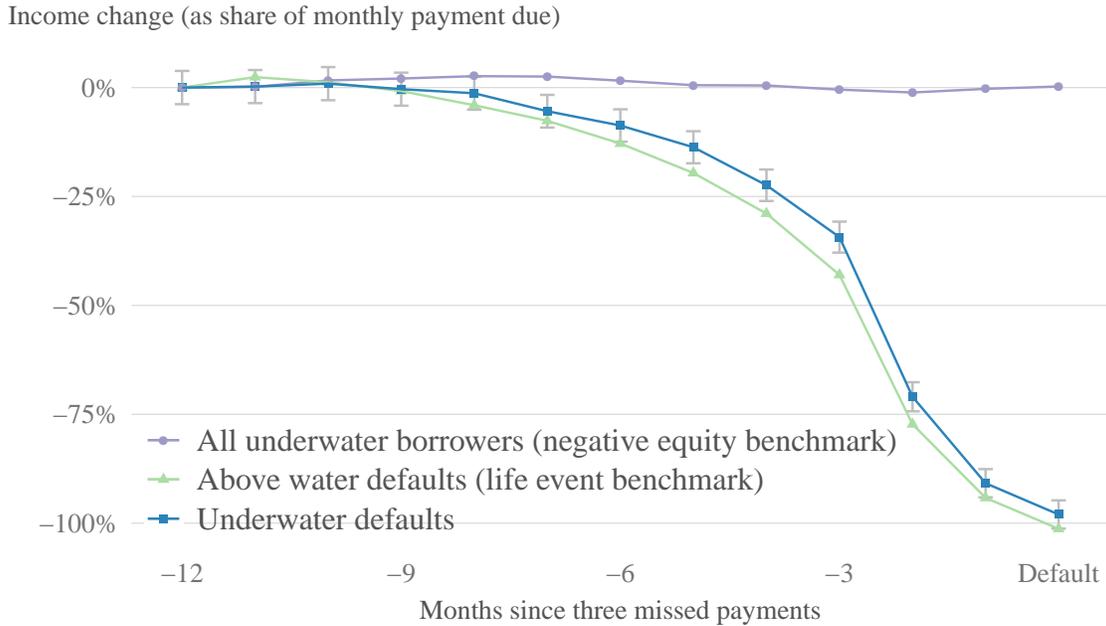
- Christie, Les.** 2010. “How foreclosure impacts your credit score.” *CNN Money*.
- Corbae, Dean, and Erwan Quintin.** 2015. “Leverage and the Foreclosure Crisis.” *Journal of Political Economy*, 123(1): 1–65.
- Cunningham, Chris, Kristopher Gerardi, and Lily Shen.** 2020. “The Double Trigger for Mortgage Default: Evidence from the Fracking Boom.” *Management Science*, 1–22.
- Dávila, Eduardo.** 2016. “Using elasticities to derive optimal bankruptcy exemptions.” *The Review of Economic Studies*, 87(2): 780–913.
- DeFusco, Anthony, and John Mondragon.** 2020. “No Job, No Money, No Refi: Frictions to Refinancing in a Recession.” *The Journal of Finance*, 75(5): 2327–2376.
- Deng, Yongheng, John M. Quigley, and Robert Van Order.** 2000. “Mortgage Terminations, Heterogeneity and the Exercise of Mortgage Options.” *Econometrica*, 68(2): 275–307.
- Diamond, Rebecca, Adam Guren, and Rose Tan.** 2020. “The Effect of Foreclosures on Homeowners, Tenants, and Landlords.” National Bureau of Economic Research Working Paper 27358.
- Diamond, William, and Tim Landvoigt.** 2019. “Credit Cycles with Market Based Household Leverage.” Social Science Research Network SSRN Scholarly Paper ID 3318481, Rochester, NY.
- Dobkin, Carlos, Amy Finkelstein, Raymond Kluender, and Matthew J. Notowidigdo.** 2018. “The Economic Consequences of Hospital Admissions.” *American Economic Review*, 108(2): 308–352.
- Elul, Ronel, Aaron Payne, and Sebastian Tilson.** 2021. “Owner-Occupancy Fraud and Mortgage Performance.” Federal Reserve Bank of Philadelphia Working Paper 19-53.
- Epperson, James F., James B. Kau, Donald C. Keenan, and Walter J. Muller.** 1985. “Pricing Default Risk in Mortgages.” *Real Estate Economics*, 13(3): 261–272.
- Experian, and Oliver Wyman.** 2009. “Experian–Oliver Wyman market intelligence report: understanding strategic default in mortgages. Part I.”
- Fagereng, Andreas, Magne Mogstad, and Marte Ronning.** 2021. “Why Do Wealthy Parents Have Wealthy Children?” *Journal of Political Economy*, 129(3): 703–756.
- Fannie Mae.** 2011. “Selling Guide.”
- Fannie Mae.** 2020. “Servicing Guide.”
- Farrell, Diana, Kanav Bhagat, and Chen Zhao.** 2018. “Falling Behind: Bank Data on the Role of Income and Savings in Mortgage Default.” JPMorgan Chase Institute.
- Farrell, Diana, Kanav Bhagat, Peter Ganong, and Pascal Noel.** 2017. “Mortgage Modifications after the Great Recession: New Evidence and Implications for Policy.” JPMorgan Chase Institute.
- Foote, Christopher L., and Paul S. Willen.** 2018. “Mortgage-Default Research and the Recent Foreclosure Crisis.” *Annual Review of Financial Economics*, 10(1): 59–100.
- Foote, Christopher L., Kristopher Gerardi, and Paul S. Willen.** 2008. “Negative equity and foreclosure: Theory and evidence.” *Journal of Urban Economics*, 64(2): 234–245.
- Foster, Chester, and Robert Van Order.** 1984. “An Option-Based Model of Mortgage Default.” *Housing Finance Review*, 3: 351–372.

- Ganong, Peter, and Pascal Noel.** 2019. “Consumer Spending during Unemployment: Positive and Normative Implications.” *American Economic Review*, 109(7): 2383–2424.
- Ganong, Peter, and Pascal Noel.** 2020. “Liquidity vs. Wealth in Household Debt Obligations: Evidence from Housing Policy in the Great Recession.” *American Economic Review*, 110(10): 3100–3138.
- Garriga, Carlos, and Aaron Hedlund.** 2020. “Mortgage Debt, Consumption, and Illiquid Housing Markets in the Great Recession.” *American Economic Review*, 110(6): 1603–1634.
- Gerardi, Kristopher, Kyle F. Herkenhoff, Lee E. Ohanian, and Paul S. Willen.** 2015. “Can’t Pay or Won’t Pay? Unemployment, Negative Equity, and Strategic Default.” National Bureau of Economic Research Working Paper 21630.
- Gerardi, Kristopher, Kyle F. Herkenhoff, Lee E. Ohanian, and Paul S. Willen.** 2018. “Can’t Pay or Won’t Pay? Unemployment, Negative Equity, and Strategic Default.” *The Review of Financial Studies*, 31(3): 1098–1131.
- Ghent, Andra C., and Marianna Kudlyak.** 2011. “Recourse and Residential Mortgage Default: Evidence from US States.” *The Review of Financial Studies*, 24(9): 3139–3186.
- Giaconetti, Marco.** Forthcoming. “Idiosyncratic Risk in Housing Markets.” *The Review of Financial Studies*.
- Gilbukh, Sonia, and Paul Goldsmith-Pinkham.** 2021. “Heterogeneous Real Estate Agents and the Housing Cycle.” Mimeo.
- Goldberger, Arthur S.** 1984. “Reverse Regression and Salary Discrimination.” *The Journal of Human Resources*, 19(3): 293–318.
- Goldberg, Lawrence, and Charles A. Capone.** 1998. “Multifamily Mortgage Credit Risk: Lessons From Recent History.” *Cityscape*, 4(1): 93–113.
- Goodman, Laurie S., Roger Ashworth, Brian Landy, and Ke Yin.** 2010. “Negative Equity Trumps Unemployment in Predicting Defaults.” *The Journal of Fixed Income*, 19(4): 67–72.
- Gopalan, Radhakrishnan, Barton H. Hamilton, Ankit Kalda, and David Sovich.** Forthcoming. “Home Equity and Labor Income: The Role of Constrained Mobility.” *The Review of Financial Studies*.
- Greenwald, Daniel, Tim Landvoigt, and Stijn Van Nieuwerburgh.** 2021. “Financial Fragility with SAM?” *The Journal of Finance*, 76(2): 651–706.
- Gross, Tal, and Matthew J. Notowidigdo.** 2011. “Health insurance and the consumer bankruptcy decision: Evidence from expansions of Medicaid.” *Journal of Public Economics*, 95(7): 767–778.
- Gross, Tal, Raymond Kluender, Feng Liu, Matthew Notowidigdo, and Jialan Wang.** 2019. “Economic Consequences of Bankruptcy Reform.” Mimeo.
- Guiso, Luigi, Paola Sapienza, and Luigi Zingales.** 2013. “The Determinants of Attitudes toward Strategic Default on Mortgages.” *Journal of Finance*, 68(4): 1473–1515.
- Gupta, Arpit, and Christopher Hansman.** 2021. “Selection, Leverage and Default in the Mortgage Market.” Working Paper.
- Gupta, Arpit, Edward R. Morrison, Catherine Fedorenko, and Scott Ramsey.** 2017. “Leverage, Default, and Mortality: Evidence from Cancer Diagnoses.” Social Science Research Network SSRN Scholarly Paper ID 2583975, Rochester, NY.

- Guren, Adam M.** 2018. "House Price Momentum and Strategic Complementarity." *Journal of Political Economy*, 126(3): 1172–1218.
- Guren, Adam M, Alisdair McKay, Emi Nakamura, and Jón Steinsson.** 2021. "Housing Wealth Effects: The Long View." *The Review of Economic Studies*, 88(2): 669–707.
- Guren, Adam M, Arvind Krishnamurthy, and Timothy J McQuade.** 2019. "Mortgage Design in an Equilibrium Model of the Housing Market." National Bureau of Economic Research Working Paper 24446.
- Gyourko, Joseph, and Joseph Tracy.** 2014. "Reconciling theory and empirics on the role of unemployment in mortgage default." *Journal of Urban Economics*, 80(C): 87–96.
- Hembre, Erik.** 2018. "HAMP, Home Attachment, and Mortgage Default." Working Paper.
- Himmelstein, David U., Elizabeth Warren, Deborah Thorne, and Steffie Woolhandler.** 2005. "Illness and injury as contributors to bankruptcy." *Health Affairs (Project Hope)*, Suppl Web Exclusives: W5–63–W5–73.
- Hsu, Joanne W., David A. Matsa, and Brian T. Melzer.** 2018. "Unemployment Insurance as a Housing Market Stabilizer." *American Economic Review*, 108(1): 49–81.
- Imbens, Guido, and Jeffrey M. Wooldridge.** 2007. "Control Function and Related Methods." What's New in Econometrics.
- Indarte, Sasha.** 2019. "The Impact of Debt Relief Generosity and Liquid Wealth on Household Bankruptcy." Social Science Research Network SSRN Scholarly Paper ID 3378669, Rochester, NY.
- Kaplan, Greg, Kurt Mitman, and Giovanni L Violante.** 2020. "The Housing Boom and Bust: Model Meets Evidence." *Journal of Political Economy*, 128(9): 3285–3345.
- Keys, Benjamin J., Tomasz Piskorski, Amit Seru, and Vikrant Vig.** 2012. "Mortgage Financing in the Housing Boom and Bust." In *Housing and the Financial Crisis*. 143–204.
- Kotova, Nadia, and Anthony L. Zhang.** 2020. "Liquidity in Residential Real Estate Markets." Working Paper.
- Laufer, Steven.** 2018. "Equity Extraction and Mortgage Default." *Review of Economic Dynamics*, 28: 1–33.
- Low, David.** 2018. "Mortgage Default with Positive Equity." Working Paper.
- Mayer, Christopher, Edward Morrison, Tomasz Piskorski, and Arpit Gupta.** 2014. "Mortgage Modification and Strategic Behavior: Evidence from a Legal Settlement with Countrywide." *American Economic Review*, 104(9): 2830–2857.
- Mitman, Kurt.** 2016. "Macroeconomic Effects of Bankruptcy and Foreclosure Policies." *American Economic Review*, 106(8): 2219–2255.
- Office of the Comptroller of the Currency.** 2005. "CREDIT RISK MANAGEMENT GUIDANCE FOR HOME EQUITY LENDING."
- O'Malley, Terry.** 2021. "The Impact of Repossession Risk on Mortgage Default." *The Journal of Finance*, 76(2): 623–650.
- Palmer, Christopher.** 2015. "Why Did So Many Subprime Borrowers Default During the Crisis: Loose Credit or Plummeting Prices?" Working Paper.

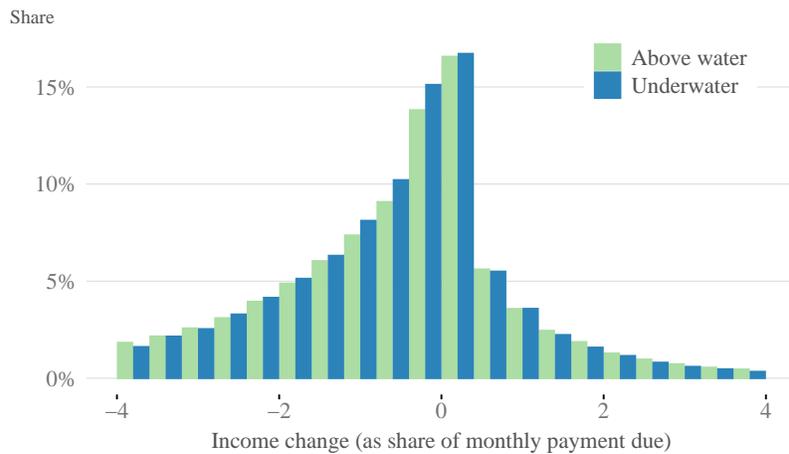
- Passero, William, Thesia I. Garner, and Clinton McCully.** 2014. "Understanding the Relationship: CE Survey and PCE." In *Improving the Measurement of Consumer Expenditures*. 181–203. University of Chicago Press.
- Pearl, Judea.** 1999. "Probabilities Of Causation: Three Counterfactual Interpretations And Their Identification." *Synthese*, 121(1): 93–149.
- Piskorski, Tomasz, Amit Seru, and James Witkin.** 2015. "Asset Quality Misrepresentation by Financial Intermediaries: Evidence from the RMBS Market." *The Journal of Finance*, 70(6): 2635–2678.
- Riddiough, Timothy J.** 1991. "Equilibrium mortgage default pricing with non-optimal borrower behavior." PhD diss. Univseristy of Wisconsin.
- Riddiough, Timothy J., and Steve B. Wyatt.** 1994. "Strategic default, workout, and commercial mortgage valuation." *The Journal of Real Estate Finance and Economics*, 9(1): 5–22.
- Rosenbaum, Paul R.** 2001. "Effects Attributable to Treatment: Inference in Experiments and Observational Studies with a Discrete Pivot." *Biometrika*, 88(1): 219–231.
- Roubini, Nouriel.** 2008. "The Forthcoming "Jingle Mail" Tsunami: 10 to 15 Million Households Likely to Walk Away from their Homes/Mortgages Leading to a Systemic Banking Crisis."
- Rubin, Donald B.** 1974. "Estimating causal effects of treatments in randomized and nonrandomized studies." *Journal of Educational Psychology*, 66(5): 688–701.
- Scharlemann, Therese, and Stephen H. Shore.** 2019. "The Effect of Mortgage Payment Size on Default and Prepayment: Evidence from HAMP Resets." *Working Paper*.
- Schelke, Thomas.** 2018. "Mortgage Default during the U.S. Mortgage Crisis." *Journal of Money, Credit and Banking*, 50(6): 1101–1137.
- Tirupattur, V, O Chang, and J Egan.** 2010. "Understanding strategic defaults." Morgan Stanley Research.
- U.S. Department of Housing and Urban Development.** 2019. "FHA Single Family Housing Policy Handbook."
- Vandell, Kerry D.** 1995. "How Ruthless Is Mortgage Default? A Review and Synthesis of the Evidence." *Journal of Housing Research*, 6(2): 245–264.
- Wooldridge, Jeffrey M.** 2010. *Econometric Analysis of Cross Section and Panel Data*. Vol. 1 of *MIT Press Books*, The MIT Press.
- Yamamoto, Tepei.** 2012. "Understanding the Past: Statistical Analysis of Causal Attribution." *American Journal of Political Science*, 56(1): 237–256.

Figure 1: Monthly Evolution of Income in Year Prior to Mortgage Default



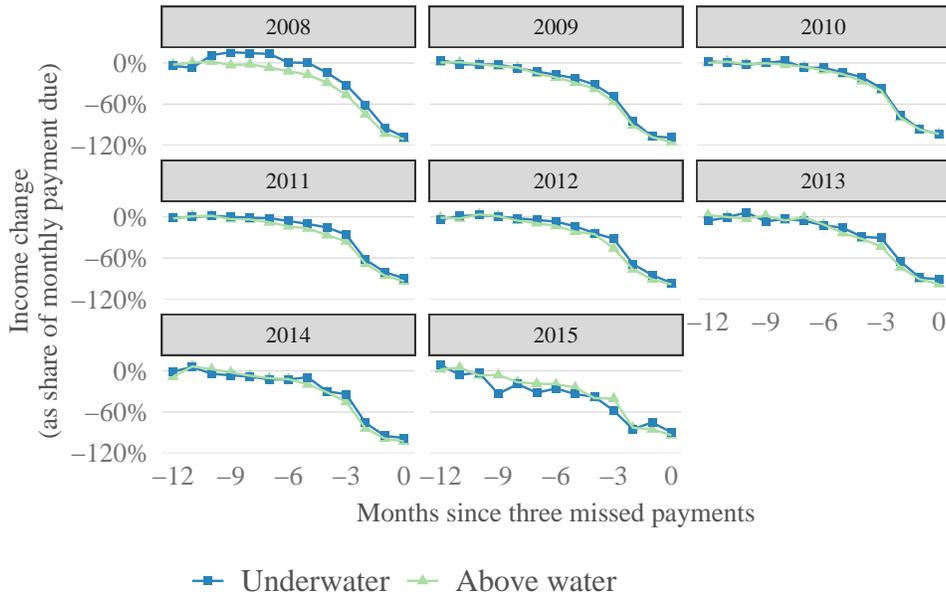
Notes: This figure describes the evolution of income in the year prior to mortgage default in the JPMCI data. The squares show mean income of underwater borrowers in comparison to two benchmarks: income for all underwater borrowers in circles, which captures the negative equity (strategic) benchmark, and income for above water borrowers in triangles, which captures the negative life event benchmark. The gray error bars show the 95 percent confidence interval for the difference between the underwater defaulters and the above water defaulters, centered at the average income drop for the underwater defaulters. Income is normalized by average payment due, which is computed separately for above and underwater borrowers over the first three months of the series. Default is defined as three missed mortgage payments. See Section 4 for details.

Figure 2: Distribution of Income Change in Year Prior to Mortgage Default



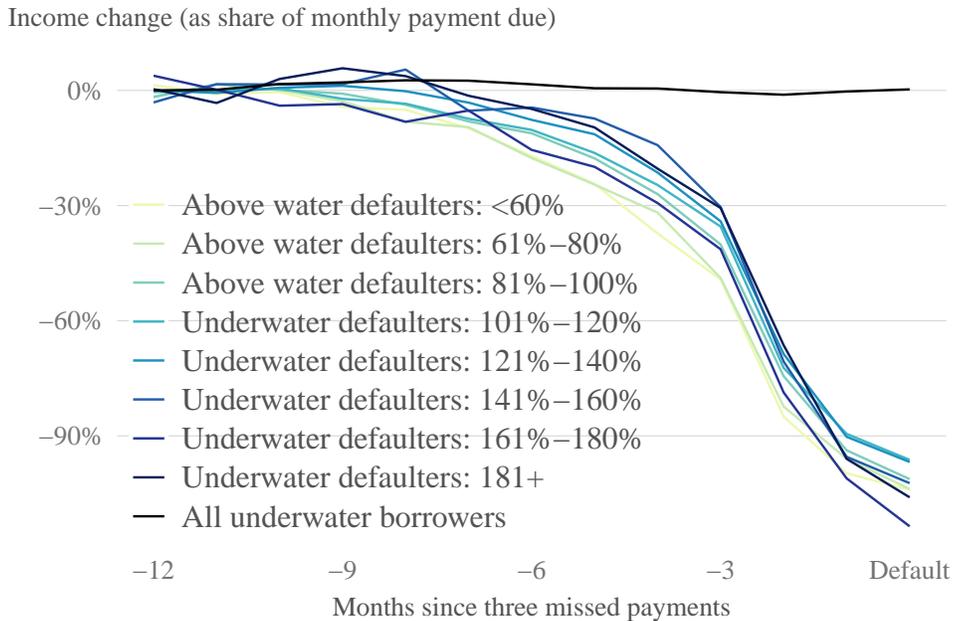
Notes: This figure shows the full distribution of the income change in Figure 1 from the three months around default relative to the first three months of the series (as in equation 7) and is truncated at -4 and 4 to improve readability.

Figure 3: Heterogeneity in Income Drop Before Default by Year



Notes: This figure shows the evolution of income in the year prior to mortgage default in the JPMCI data separately for each year between 2008 and 2015. Default is defined as three missed mortgage payments. Mean income is normalized by average payment due, which is computed separately by LTV and year bin. Monthly changes are reported relative to the average in the first three months of the series. See Section 4.3 for details.

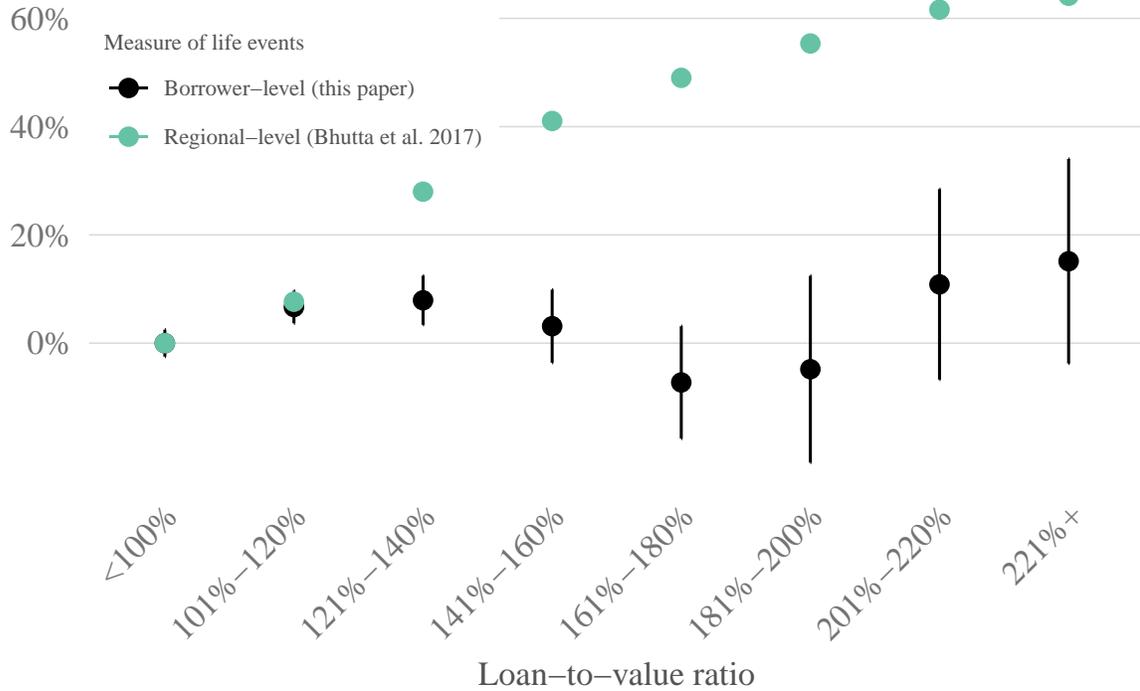
Figure 4: Heterogeneity in Income Drop by Loan-to-Value Ratio



Notes: This figure shows the evolution of income in the year prior to mortgage default in the JPMCI data separately by granular LTV bins. Mean income is normalized by average payment due, which is computed separately by LTV bin. Monthly changes are reported relative to the average in the first three months of the series. See Section 4.3 for details.

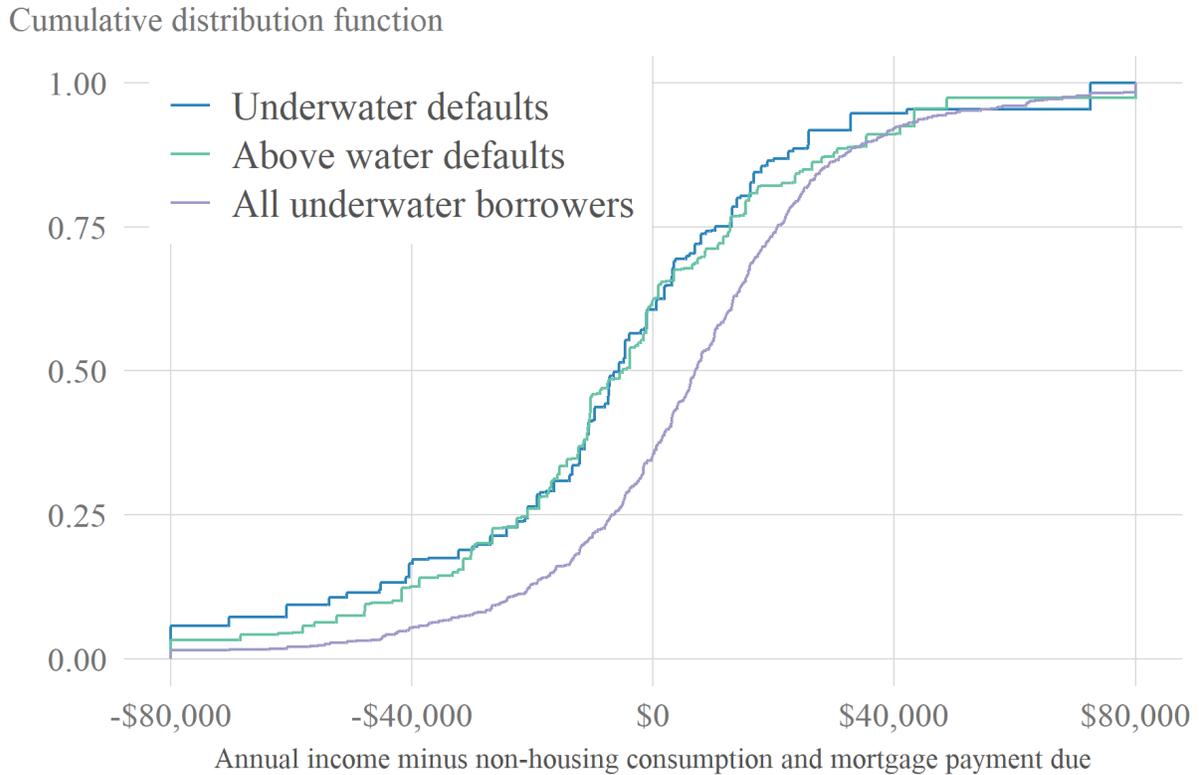
Figure 5: Share of Strategic Defaults by Loan-to-Value Ratio

Share of strategic defaults



Notes: This figure compares estimates of the share of strategic defaults using regional measures of life events (from Bhutta, Dokko and Shan 2017) and individual measures of life events in the JPMCI bank account data. The bank account estimates report $1 - \hat{\alpha}_{\text{life event}}$ using equation (4). The Bhutta, Dokko and Shan (2017) estimates come from Figure 6 of that paper, where the proportion of strategic defaults is the difference between the predicted default rate and the predicted default rate ignoring the equity effect at each LTV bin. The vertical lines show 95 percent confidence intervals for our estimates.

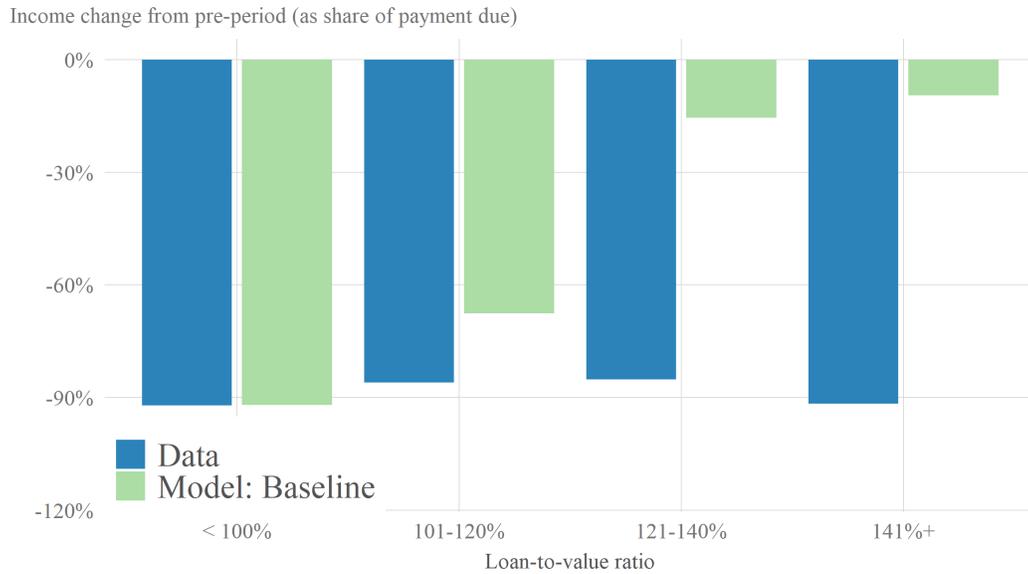
Figure 6: Distribution of Available Resources is Same for Above Water and Underwater Defaulters in PSID



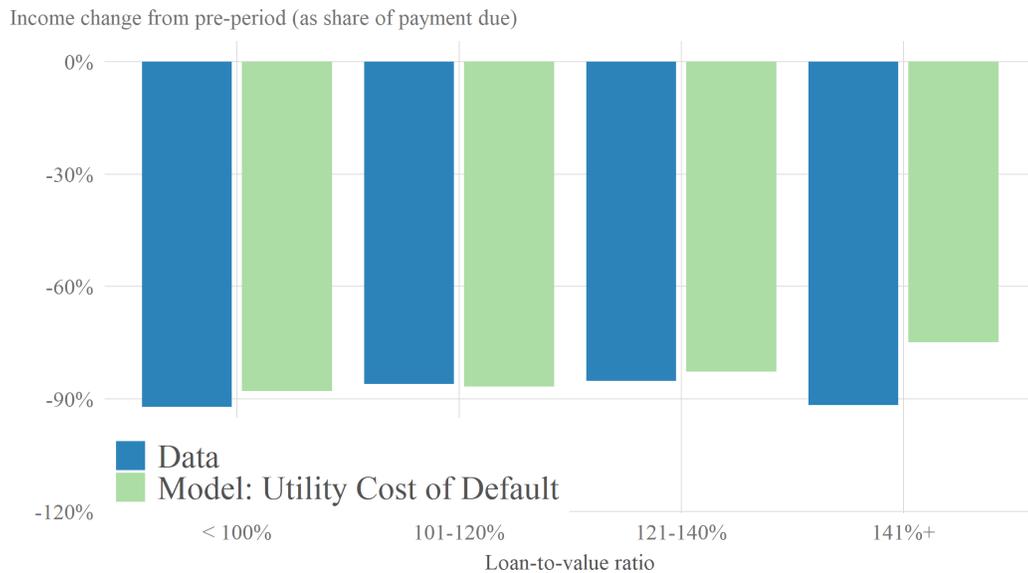
Notes: This figure shows the cumulative distribution function of available resources by home equity and default status in the Panel Study of Income Dynamics (PSID). We follow Gerardi et al. (2018) in defining available resources as annual income minus non-housing consumption and mortgage payment due. We winsorize this variable at +/- \$80,000. The figure shows the distribution separately for above water defaulters, underwater defaulters, and all underwater borrowers. See Section 5 for details.

Figure 7: Income Drop at Default Compared to Prior Theoretical Literature

(a) Baseline Model



(b) Model with Estimated Utility Cost of Default



Notes: This figure compares income drop at default by home equity in the Campbell and Cocco (2015) structural model and the JPMCI data. The model bars in panel (a) show results from that paper’s baseline model where a borrower defaults when the utility from renting is greater than the utility from paying a mortgage. The model bars in panel (b) show a model variant with a utility cost of mortgage default which is equal to a 25 percent decrease in the constant-equivalent consumption stream. The y-axis is the change in annual income divided by the annual mortgage payment due in the model and the change in monthly income from one year prior divided by the monthly mortgage payment due in the data. The x-axis is the loan-to-value ratio in the year of default. See Section 6 for details.

Table 2: Potential Outcomes Model of the Mortgage Default Decision

Type	Decision rule	$Y(\underbrace{T^*}_{\text{life event}}, \underbrace{G}_{\text{negative equity}})$			
		$Y(0,0)$	$Y(1,0)$	$Y(0,1)$	$Y(1,1)$
Strategic (ST)	Negative equity is necessary and sufficient	0	0	1	1
Cash-flow (CF)	Negative life event is necessary and sufficient	0	1	0	1
Double trigger (DT)	Both negative equity and negative life event are necessary	0	0	0	1

Notes: This table shows the different combination of potential outcomes from the environment described in Section 2. The fact that there are three potential outcome types arises from Assumptions 1 and 4 and our focus on defaulters who are already underwater. See Section 2.3.2 for details.

Table 3: Summary Statistics by Home Equity

Variable	Above water	Underwater
Combined loan-to-value ratio (%)	75	125
Bank account income (\$)	4,279	4,541
Bank account balance (\$)	1,523	1,766
Property value (\$)	248,471	222,148
Monthly mortgage payment due (\$)	1,143	1,364
Age	49	47
Share with joint deposit account	0.4	0.45
N	85,654	52,331

Notes: This table shows means describing the income sample of defaulters in the JPMCI data six months prior to mortgage default. Above water borrowers have positive home equity (combined LTV < 100) and underwater borrowers have negative home equity (combined LTV > 100). We refer to this variable as *combined* loan-to-value (LTV) ratio because it includes both first lien and second lien mortgage debt. Default is defined as three missed mortgage payments.

Table 4: Income Drop at Default by Home Equity

	<i>Dependent variable:</i>				
	Mean	Median	p25	p75	Mean (w/ME correction)
	(1)	(2)	(3)	(4)	(5)
Date of default	-0.921 (0.008)	-0.758 (0.009)	-0.567 (0.010)	-1.259 (0.016)	-0.916 (0.010)
Date of default * underwater	0.051 (0.012)	-0.039 (0.015)	-0.059 (0.016)	0.057 (0.022)	0.037 (0.018)
N mortgages	137,985	137,985	137,985	137,985	137,985
Observations	827,910	827,910	827,910	827,910	827,910

Notes: This table describes the income drop at default in the JPMCI data. The sample uses a balanced panel with a pre-period (12 to 10 months prior to default) and a default period (2 months prior to default through the month of default), so the number of observations is six times the number of mortgages. The dependent variable is the ratio of monthly income to average monthly payment due in the pre-period. The regression specification is $\frac{Income_t}{Payment_{pre}} = \lambda + \kappa(LTV > 100) + \gamma \mathbf{1}(t = -2, -1, 0) + \beta \mathbf{1}(t = -2, -1, 0) \times LTV > 100 + \varepsilon$, as shown in equation (7). The table reports estimates for the change in income during the default period ($\hat{\gamma}$) and the interaction with being underwater ($\hat{\beta}$). Column (1) analyzes the mean change in income. Columns (2), (3), and (4) show the 50th, 25th and 75th percentiles of the change in income respectively. Column (5) repeats column (1), but using the measurement error (ME) correction for LTV from equation (10). Above water borrowers have positive home equity ($LTV < 100$) and underwater borrowers have negative home equity ($LTV > 100$). Default is defined as three missed mortgage payments. Standard errors in OLS regressions are clustered by mortgage and in quantile regressions are computed by bootstrapping over mortgages.

Table 5: Share of Defaults Causally Attributable to Life Events ($\hat{\alpha}_{\text{life event}}$)

(a) Estimates for All Borrowers

Category	$\hat{\alpha}_{\text{life event}}$ (SE)
Baseline	0.944 (0.013)
Change from $t=-12$ to $t=0$	0.967 (0.016)
Foreclosure	0.963 (0.017)
LTV measurement error correction	0.959 (0.017)

(b) Heterogeneity

Category	$\hat{\alpha}_{\text{life event}}$ (SE)
Year	
2008	0.914 (0.058)
2009	0.947 (0.035)
2010	0.987 (0.029)
2011	0.942 (0.028)
2012	0.941 (0.032)
2013	0.934 (0.041)
2014	0.944 (0.054)
2015	0.955 (0.105)
LTV	
101-120	0.933 (0.015)
121-140	0.921 (0.023)
141-160	0.969 (0.034)
161-180	1.073 (0.053)
181-200	1.048 (0.088)
201-220	0.891 (0.090)
221+	0.849 (0.096)
Mortgage type	
Fixed	0.936 (0.015)
Adjustable	1.021 (0.029)
Non recourse states	1.029 (0.029)
Three consecutive missed payments	0.876 (0.013)

Notes: This table reports estimates of $\hat{\alpha}_{\text{life event}}$, which is the share of defaults causally attributable to life events, in the JPMCI data. $\hat{\alpha}_{\text{life event}}$ is constructed using equation (4). Panel (a) shows estimates for all borrowers. The first row is the baseline specification, using the regression in equation (7) and defining default as three missed payments. The second row implements equation (7) but looks at the change from the first month ($t = -12$) to the last month ($t = 0$), rather than the average of the first three months and last three months. The third row dates default as the month of foreclosure initiation. The fourth row implements the two-sample IV measurement error correction procedure for LTV described in Section 4.3. Panel (b) shows estimates for specific subsets of borrowers using the baseline specification. See Sections 4.2 and 4.3 for details.

Table 6: Impact of Negative Equity on Default

(a) Chase Sample

	(1)	(2)	(3)	(4)
Underwater	1.508 (0.065)	1.313 (0.070)	0.356 (0.065)	0.315 (0.047)
LTV fitted residuals			2.085 (0.136)	2.009 (0.109)
$\hat{\alpha}_{\text{negative equity}}$	0.779 (0.014)	0.731 (0.019)	0.299 (0.046)	0.27 (0.034)
Region-Year FEs	Y	Y	Y	N
CBSA FEs	Y	Y	Y	Y
Origination year FEs	N	N	N	Y
Borrower and loan characteristics	N	Y	Y	Y
Instrument	-	-	Cyclicality-HPI	Cyclicality-Month
First stage partial F-Stat	-	-	121.97	13.08
Log Likelihood	-424,526	-419,933	-416,127	-414,939
Observations	1,270,695	1,270,695	1,270,695	1,270,695

(b) CRISM Sample

	(1)	(2)	(3)	(4)
Underwater	1.385 (0.034)	0.947 (0.034)	0.266 (0.042)	0.242 (0.037)
LTV fitted residuals			1.640 (0.085)	1.601 (0.077)
$\hat{\alpha}_{\text{negative equity}}$	0.750 (0.008)	0.612 (0.013)	0.233 (0.026)	0.215 (0.026)
Region-Year FEs	Y	Y	Y	N
CBSA FEs	Y	Y	Y	Y
Origination year FEs	N	N	N	Y
Borrower and loan characteristics	N	Y	Y	Y
Instrument	-	-	Cyclicality-HPI	Cyclicality-Month
First stage partial F-Stat	-	-	1201.12	126.35
Log Likelihood	-455,846	-445,190	-442,536	-442,612
Observations	1,434,225	1,434,225	1,434,225	1,434,225

Notes: This table reports maximum-likelihood estimates of the default hazard model from equation (11). Panel (a) reports estimates for the main Chase analysis sample. Panel (b) reports estimates for the CRISM analysis sample. Underwater is an indicator equal to 1 when $LTV > 1$. Columns (3) and (4) instrument for underwater using the control function approach by including the LTV fitted residuals from equation (12). CBSA-level cyclicality is from Guren et al. (2021). The instrument in column (3) interacts CBSA cyclicality with the log annual change in the regional price index. The instrument in column (4) interacts CBSA cyclicality with calendar-month fixed effects. For computational feasibility we collapse to annual data when estimating the second stage (equation 11), so the number of observations reported is the number of borrower-years. Standard errors are clustered at the CBSA level. The table also reports the share of underwater defaults for which negative equity is a necessary condition, i.e. $\hat{\alpha}_{\text{negative equity}}$. Denoting the coefficient on Underwater as $\hat{\delta}$ from equation (11), then $\hat{\alpha}_{\text{negative equity}} = 1 - \exp(-\hat{\delta})$. See Section 4.5 for details.

Online Appendix to “Why Do Borrowers Default on Mortgages? A New Method for Causal Attribution”

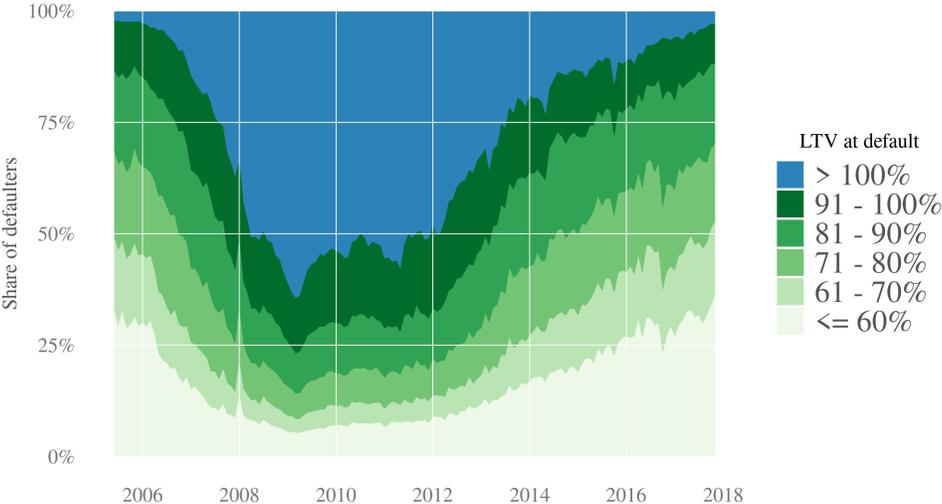
Peter Ganong and Pascal Noel

Contents

A	Figures and tables	1
B	Data appendix	30
C	Proof of proposition 1	32
D	Relaxing Assumption 2	33
E	Bias-Variance Trade-off	37
F	Income after default	40
G	Stigma cost of default	42

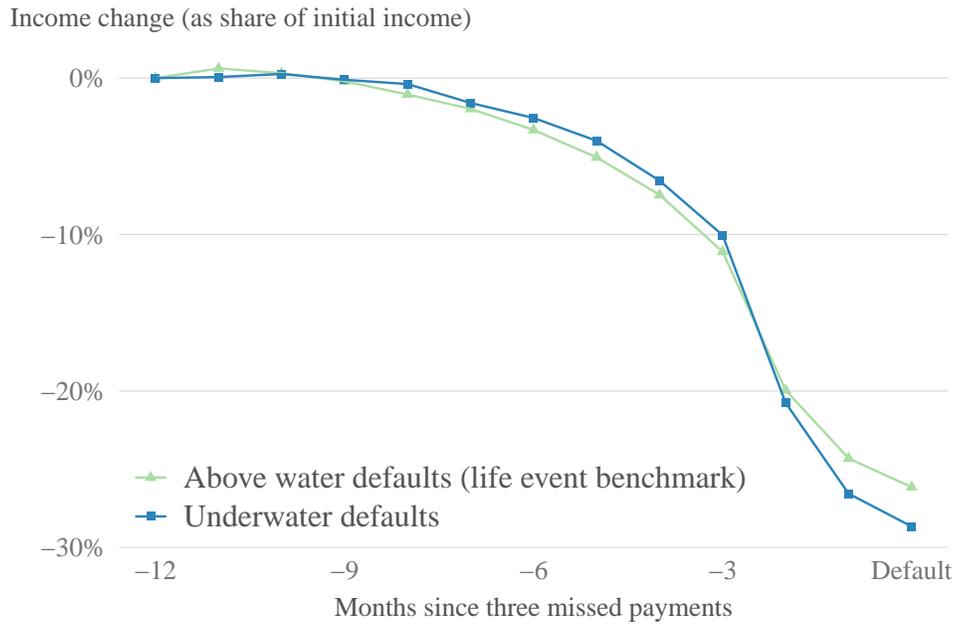
A Figures and tables

Figure A-1: Prevalence of Above Water Mortgage Default



Notes: This figure shows the distribution of the loan-to-value (LTV) ratio at default in the Credit Risk Insight Servicing McDash (CRISM) data. Default is defined as three missed payments.

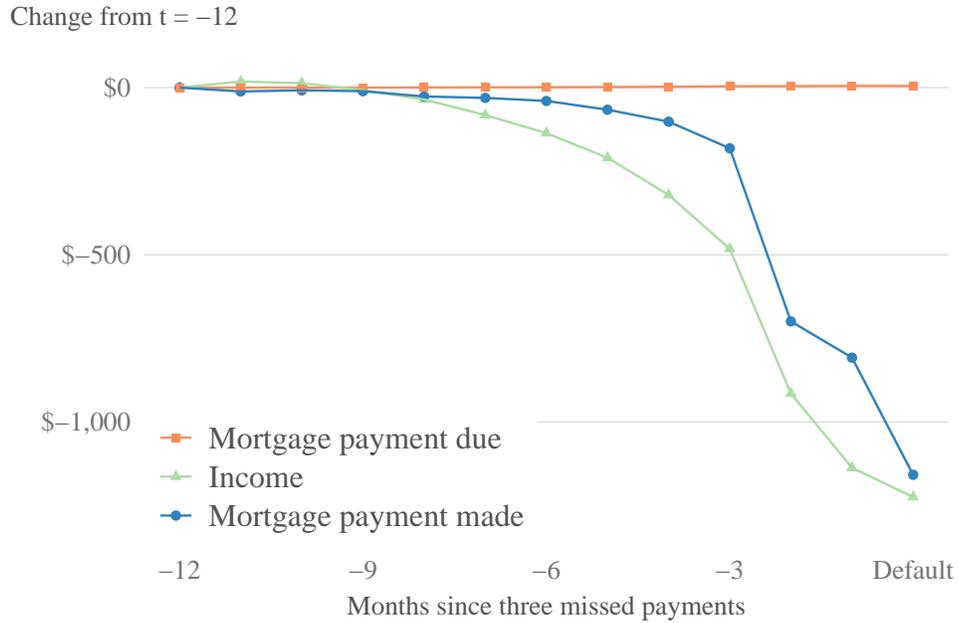
Figure A-2: Income Change as Share of Initial Income by Home Equity



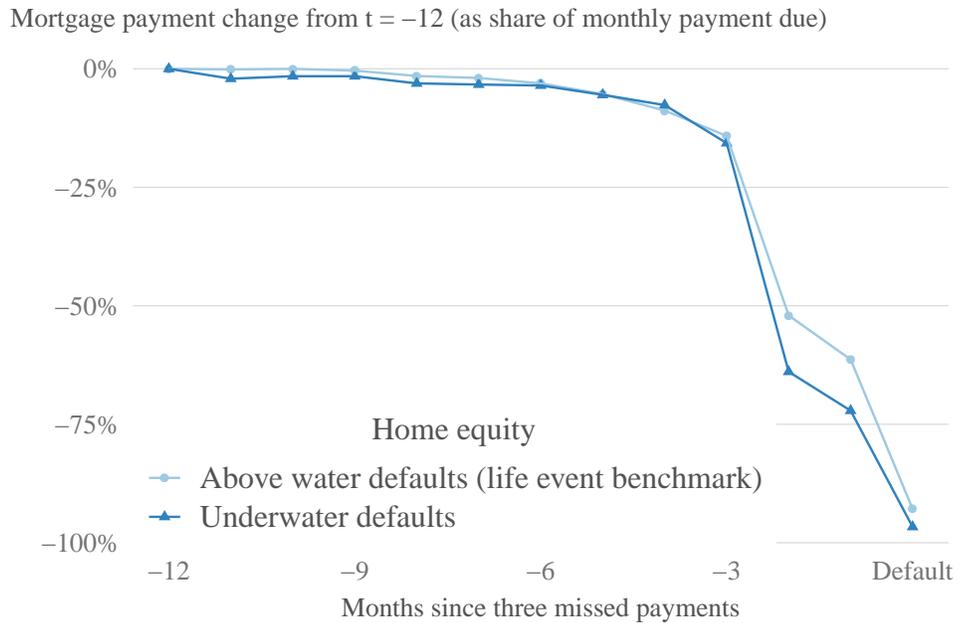
Notes: This figure replicates Figure 1 using a dependent variable of the change in monthly income divided by the average of the monthly income in months -12, -11 and -10 prior to default.

Figure A-3: Payment Due And Payment Made Prior to Default

(a) Payment Due

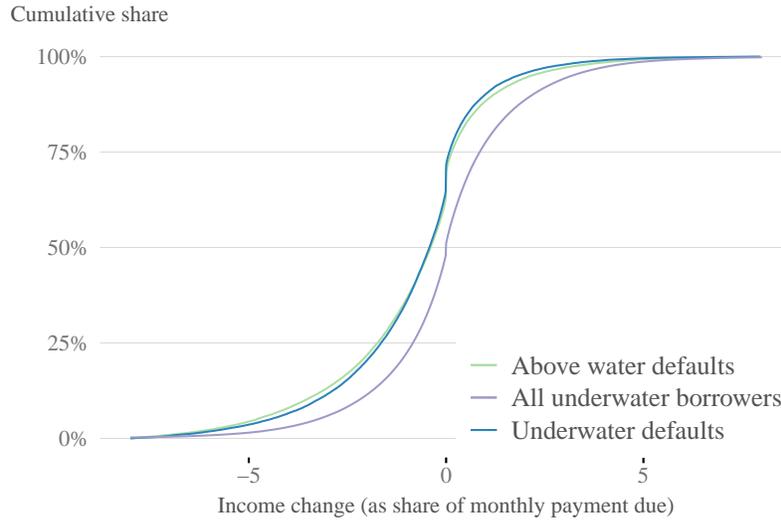


(b) Mortgage Payments by Home Equity



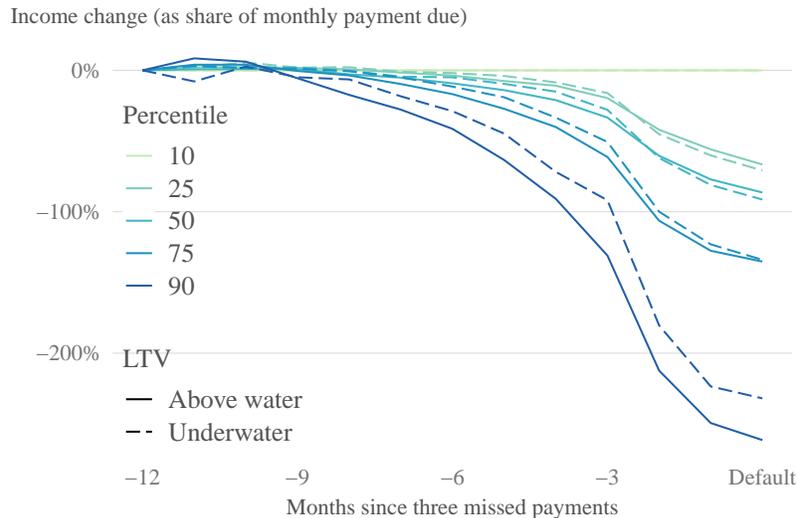
Notes: The top panel shows mortgage payment due, average income, and mortgage payment made in the year prior to default in the JPMCI data. The bottom panel shows mortgage payment made as a share of payment due in the year prior to default in the JPMCI data.

Figure A-4: Distribution of Income Change Prior to Mortgage Default



Notes: This figure shows the cumulative distribution function for the change in income, divided by average initial payment due. Average initial payment due is computed one year prior to mortgage default and is computed separately for underwater and above water borrowers. This figure provides an alternative visualization of the histogram in Figure 2, and includes the change for all underwater borrowers. The distribution of the change in income is truncated at -8 and 8 to improve readability.

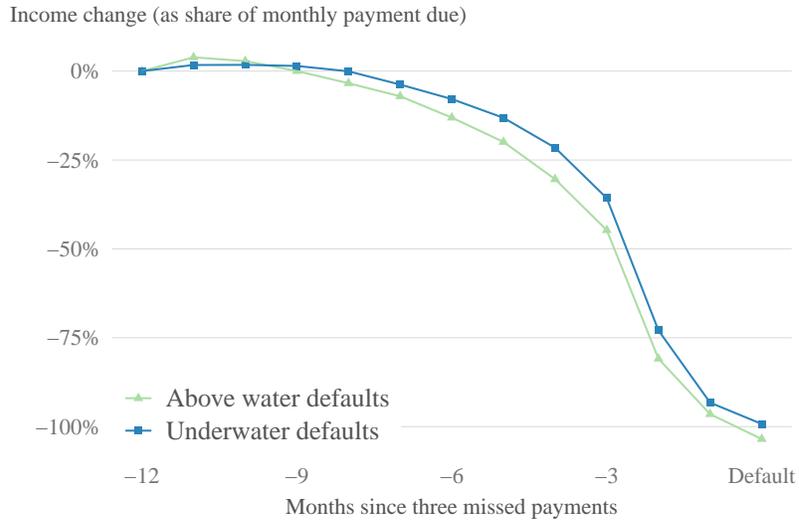
Figure A-5: Percentiles of Income Change Prior to Mortgage Default by Home Equity



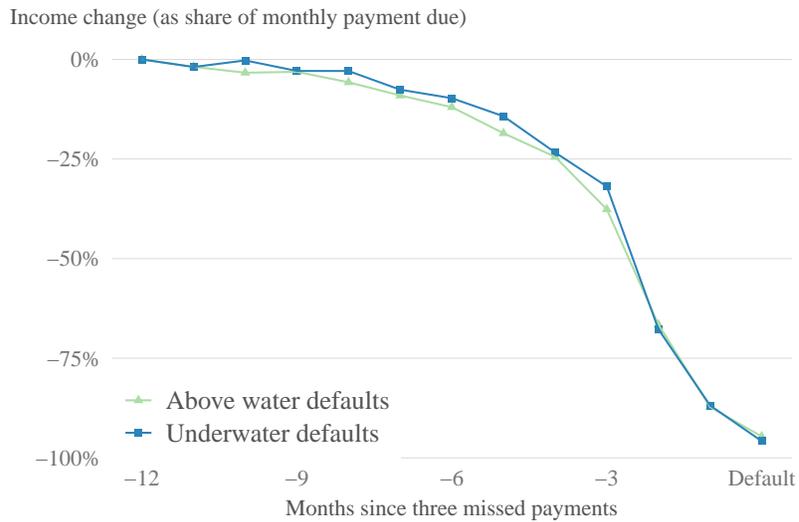
Notes: This figure shows the income change in the year prior to mortgage default for different percentiles of the income change distribution. The percentiles are calculated separately for above water and underwater borrowers. See Section 4.3 for details.

Figure A-6: Income in Year Prior to Mortgage Default by Mortgage Type and Home equity

(a) Fixed Rate Mortgages

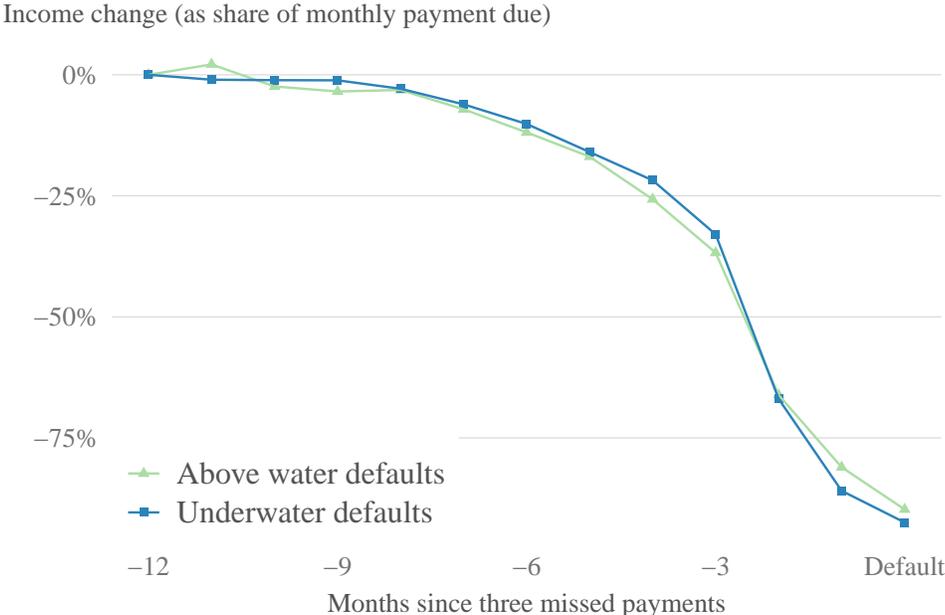


(b) Adjustable Rate Mortgages



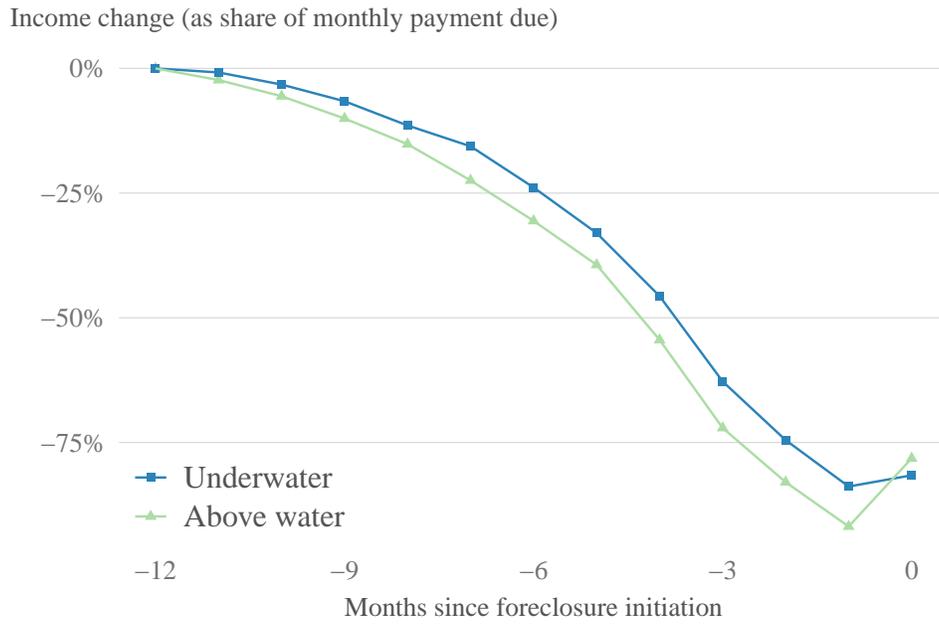
Notes: This figure replicates Figure 1 separately for fixed rate and adjustable rate mortgages.

Figure A-7: Income Prior to Default in Non-Recourse States



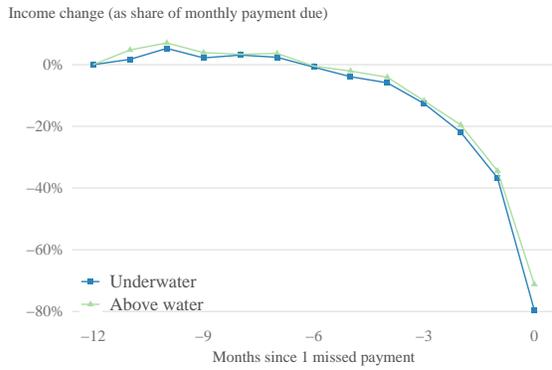
Notes: This figure replicates Figure 1 from the JPMCI data for the subset of states that do not allow mortgage lenders to sue to recover non-mortgage assets. We use the classification of non-recourse states from Ghent and Kudlyak (2011).

Figure A-8: Evolution of Income by Home Equity Prior to Foreclosure

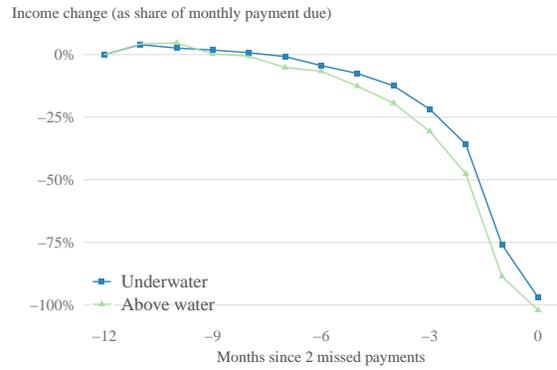


Notes: This figure replicates Figure 1 defining the date of default as the date of foreclosure initiation.

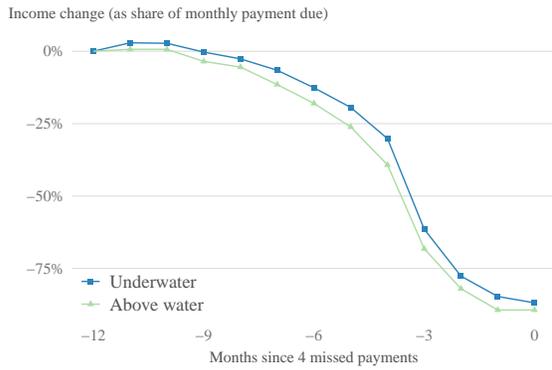
Figure A-9: Income by Alternative Missed Payment Thresholds and Home Equity



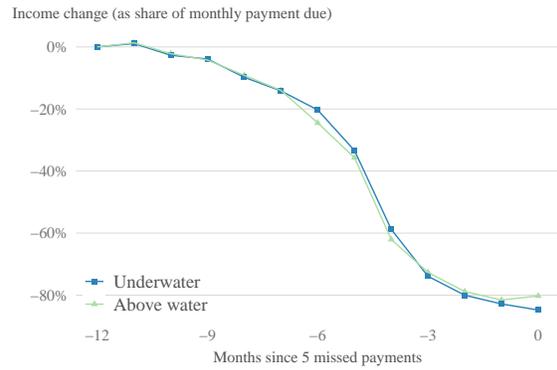
(a) One Month Past Due



(b) Two Months Past Due



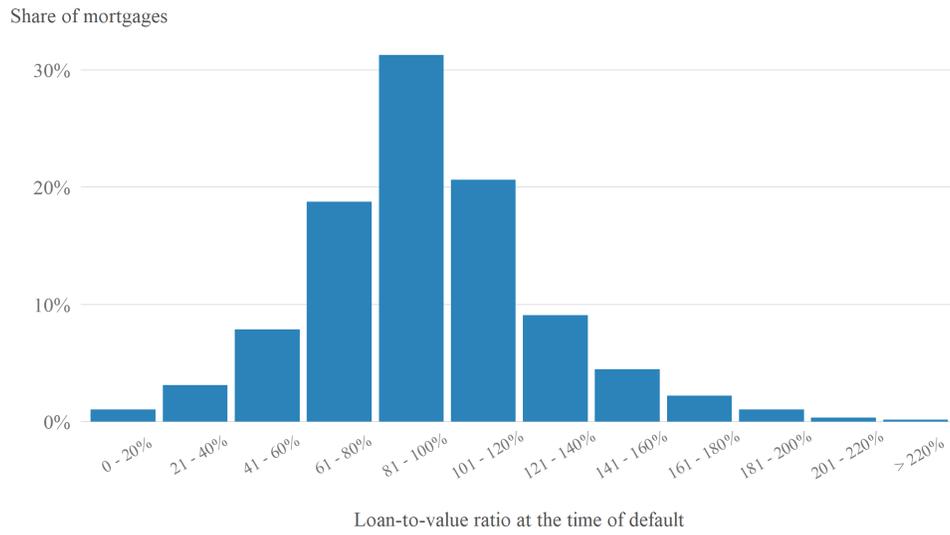
(c) Four Months Past Due



(d) Five Months Past Due

Notes: This figure replicates Figure 1 for alternative months past due thresholds.

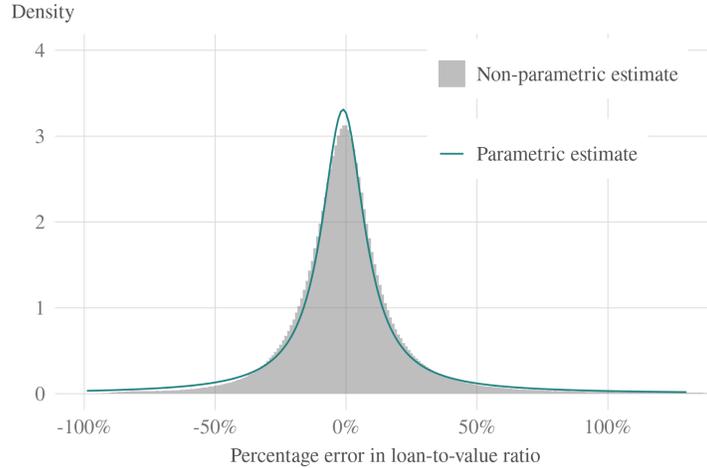
Figure A-10: Distribution of Defaulters by LTV



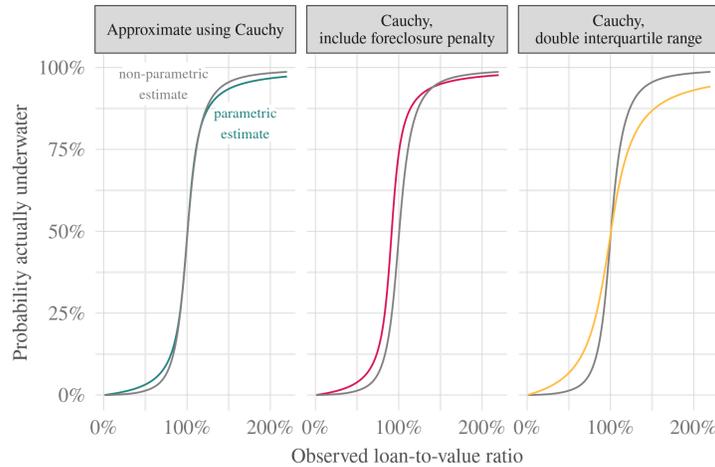
Notes: This figure shows the distribution of defaulters in the Chase analysis sample by LTV.

Figure A-11: Estimate of Measurement Error in Observed Loan-to-Value Ratio

(a) Parametric approximation



(b) Probability truly underwater by observed LTV



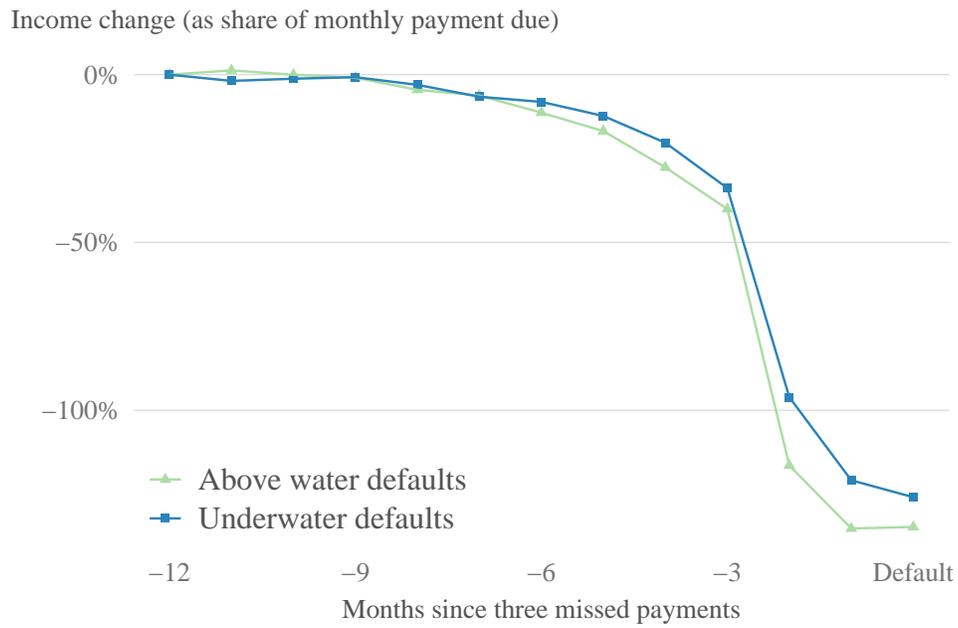
Source: Corelogic Home Price Indexes and Deed data.

Notes: This figure provides supporting analysis to the adjustment of $\hat{\alpha}_{\text{life event}}$ for measurement error in observed LTV described in Section 4.3 and shown in Table 5a. Our method for constructing the error in observed LTV largely follows Giacometti (Forthcoming). See Appendix B.3 for details.

The top panel compares the true distribution of LTV errors to a Cauchy distribution. The true distribution of errors is shown in gray bars as the inverse of the distribution of home sale price errors. Home sale price errors are $\frac{\text{PriorSalesPrice} \times \Delta \text{HomePriceIndex}}{\text{ActualSalesPrice}} - 1$. The teal line approximates the non-parametric estimate using a Cauchy distribution, which is truncated from below at -100 percent. We estimate the location and scale parameters of this distribution by minimizing the squared distance between the actual median and interquartile range and the simulated median and interquartile range.

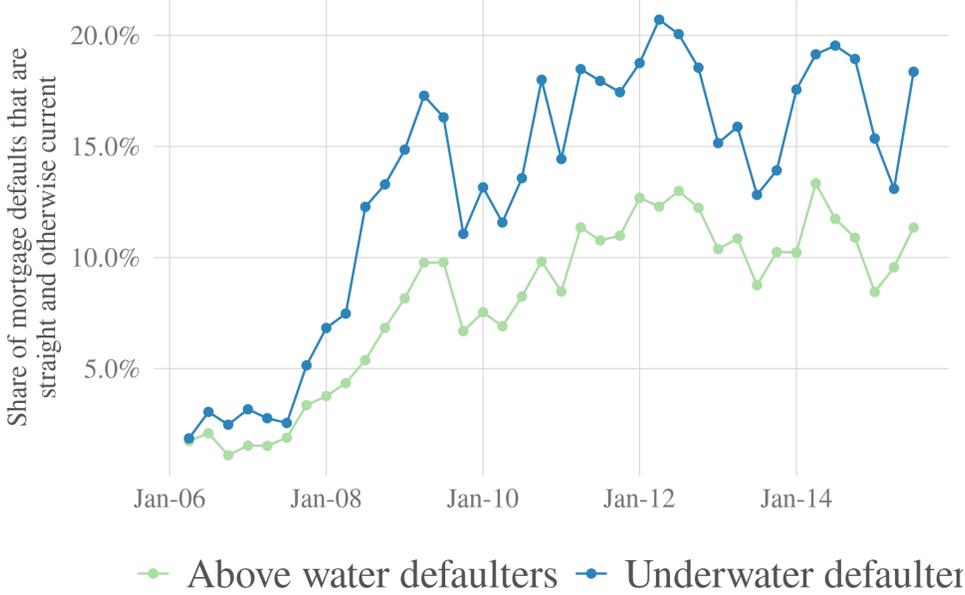
In the bottom panel, we use this parametric distribution to compute the probability that a borrower is actually underwater for a range of observed LTV values, again in teal. In the text, we refer to this function as $P(G^* = 1 | LTV)$. The bottom panel also includes two other $P(G^* = 1 | LTV)$ functions where we modify the function in teal. In the yellow line we assume that the dispersion of LTV errors (as measured by the interquartile range) is twice as large as the true range in the CoreLogic data. In the red line we assume that homes in default sell at an average discount of 9 percent relative to their observed LTV.

Figure A-12: Income Prior to Default by Consecutive Missed Payments



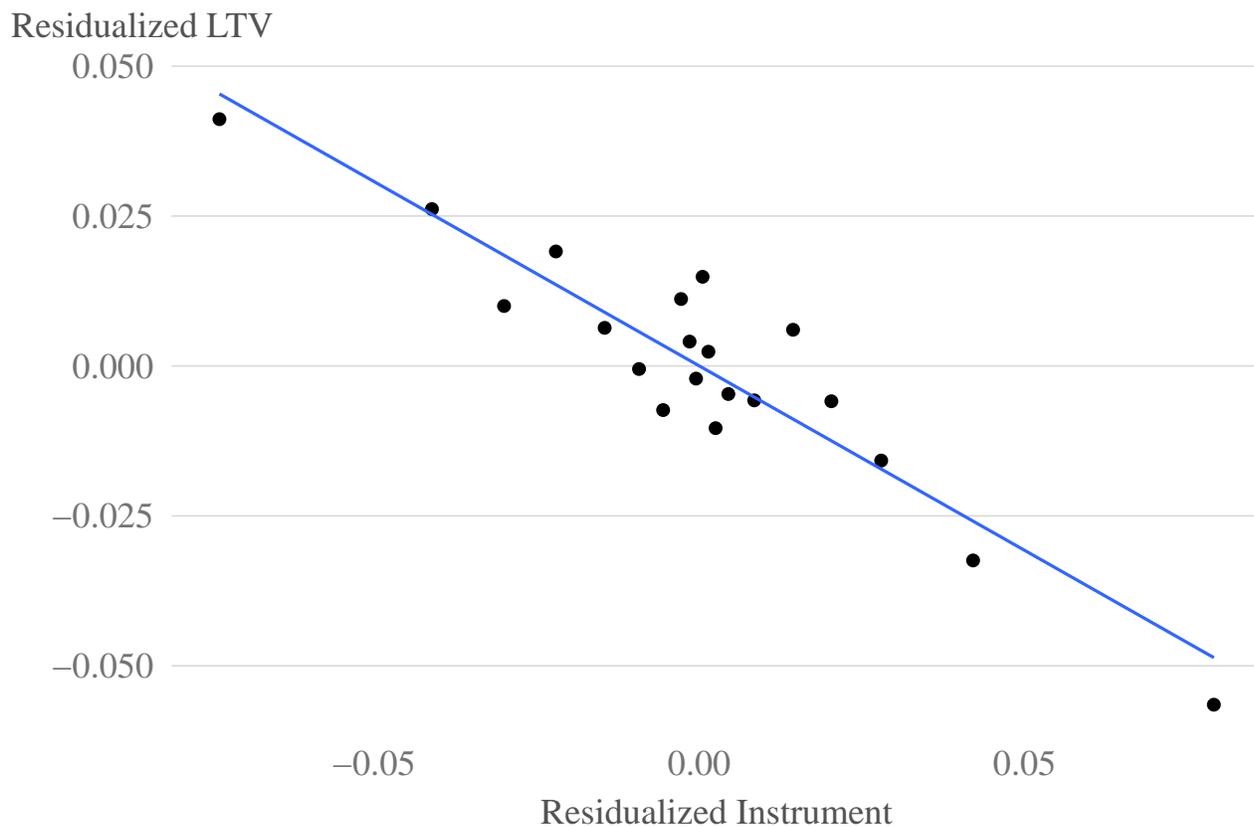
Notes: This figure replicates Figure 1 from the JPMCI data for the subset of borrowers who miss three consecutive payments. Borrowers who miss three consecutive payments are 58 percent of underwater defaults and 44 percent of above water defaults.

Figure A-13: Share of Mortgage Defaults with Consecutive Missed Payments



Notes: This figure extends the analysis in Keys et al. (2012) using the CRISM data. That paper measures the share of mortgage defaults that transition straight from 60 days past due to 180 days past due in four months, while remaining otherwise current on all non-HELOC revolving debt. We refer to such defaults as “straight and otherwise current”. The average share of defaults that meet this definition is 16.3 percent of defaults for underwater borrowers and 10 percent of defaults for above water borrowers. Thus, the excess share of straight and otherwise current defaults for underwater borrowers is 6.3 percent.

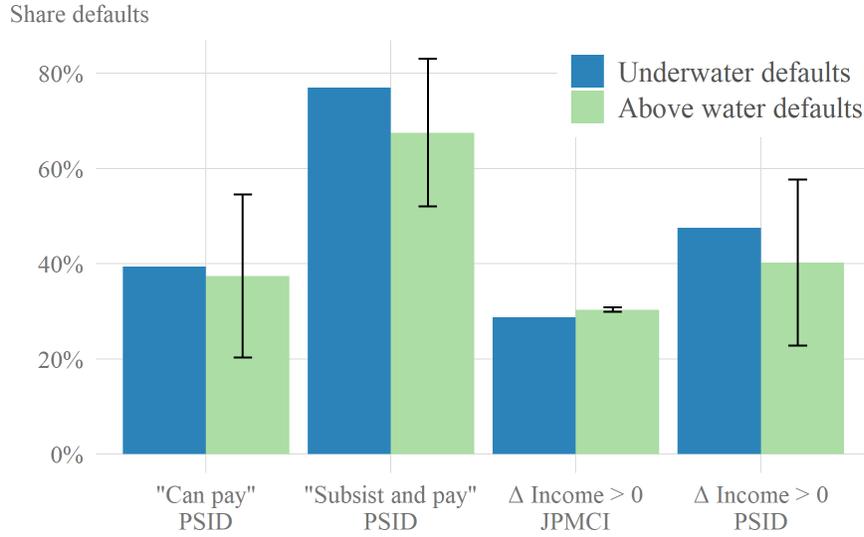
Figure A-14: Cyclical Instrument First Stage



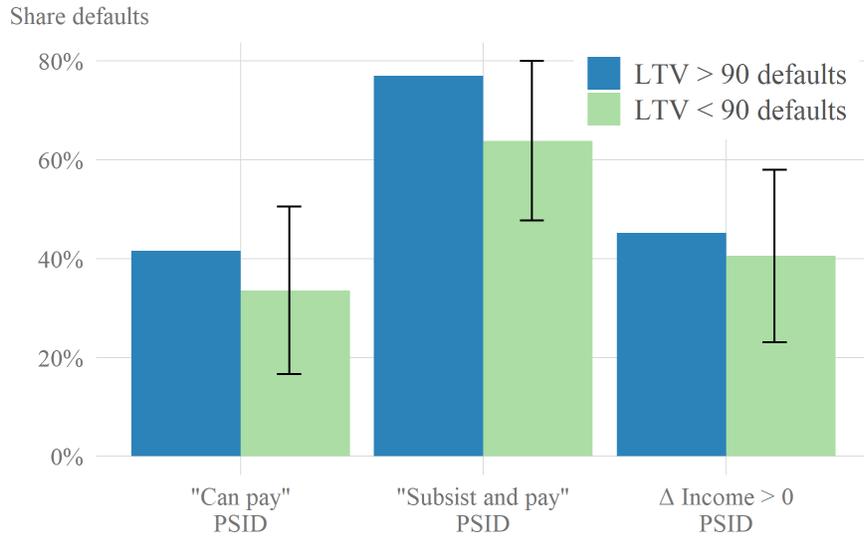
Notes: This figure shows a binned scatter plot of the first stage relationship between the cyclical instrument and LTV in the Chase sample. This corresponds to equation (12). Both the instrument and LTV are residualized against all fixed effects and controls.

Figure A-15: Alternative Measures of Strategic Default

(a) Share of Defaults



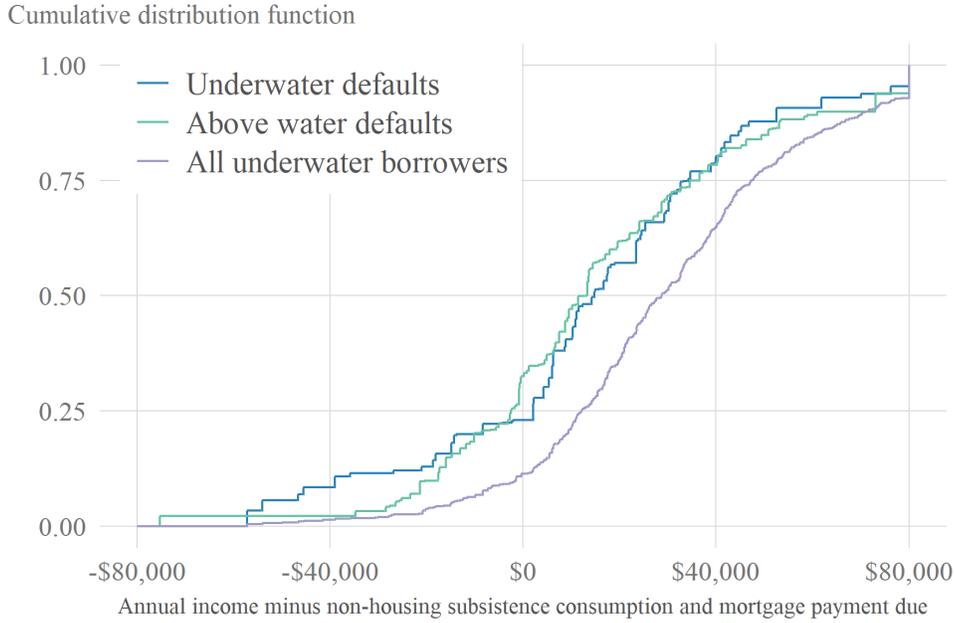
(b) Share of Defaults Using Loan-to-Value (LTV) Cutoff of 90



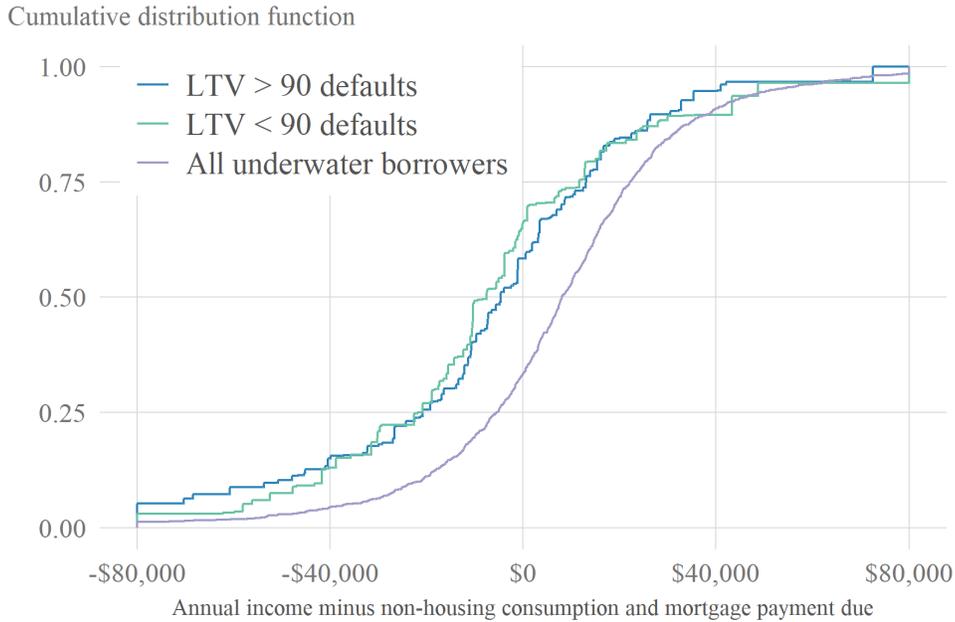
Notes: This figure compares the distribution of income by home equity and default status in the Panel Study of Income Dynamics (PSID) and the bank account data. Gerardi et al. (2018) measure mortgage affordability using income y , mortgage payment m , and non-housing consumption c . That paper classifies a borrower as *can-pay* if she can afford the mortgage without cutting consumption ($y - m - c_{predefault} > 0$) and as *subsist-and-pay* if she can afford a subsistence consumption level and pay her mortgage ($y - m - c_{subsistence} > 0$). See Section 5 for details on these definitions. Panel (a) reports the share of defaults that are classified as strategic using three different empirical criteria: can-pay, subsist-and-pay, and a positive change in income. The vertical lines indicate 95 percent confidence intervals for the difference in shares between above and underwater. Panel (b) reproduces the PSID analysis from panel (a), classifying defaults by whether the borrower's LTV is above 90, which is the LTV cutoff used in Gerardi et al. (2018).

Figure A-16: Alternative Measures of Strategic Default – Distributions

(a) Available Resources Using Subsistence Measure

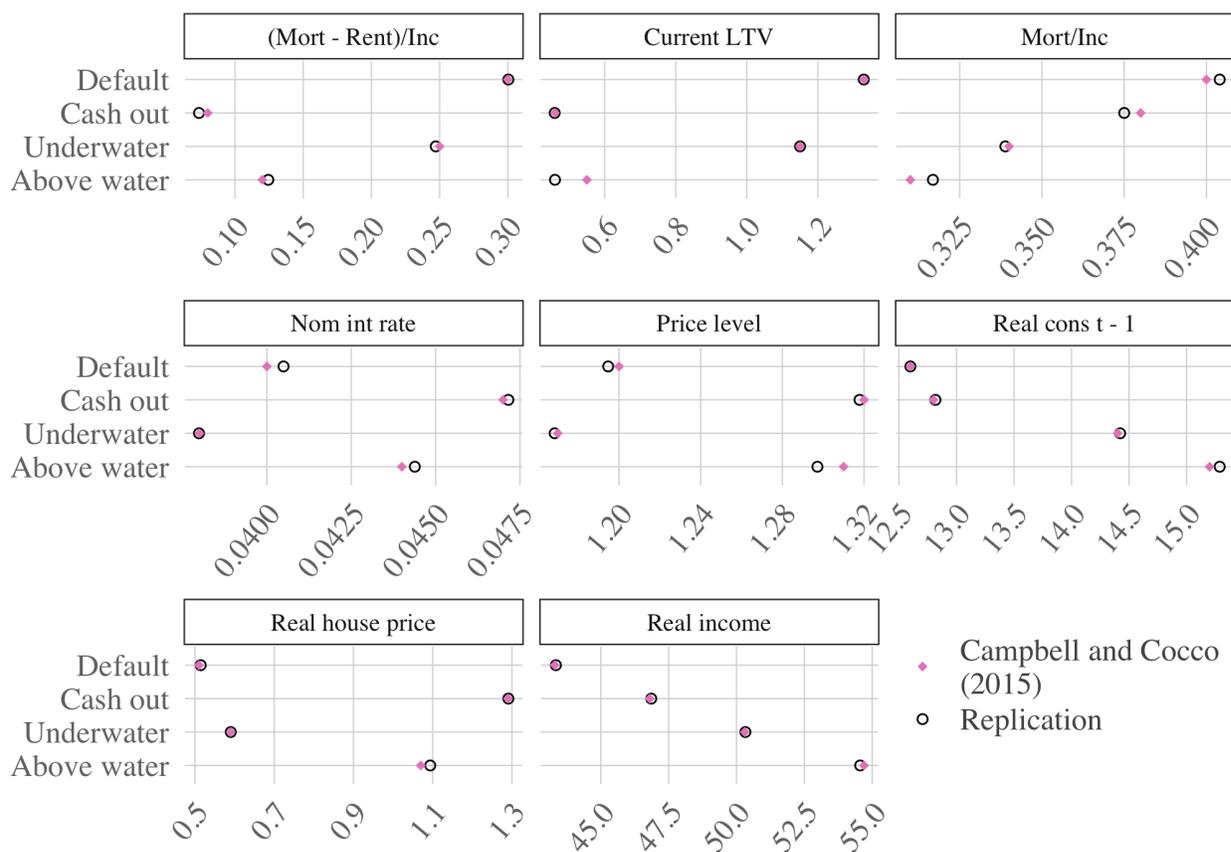


(b) Available Resources Using Loan-to-value (LTV) Cutoff of 90



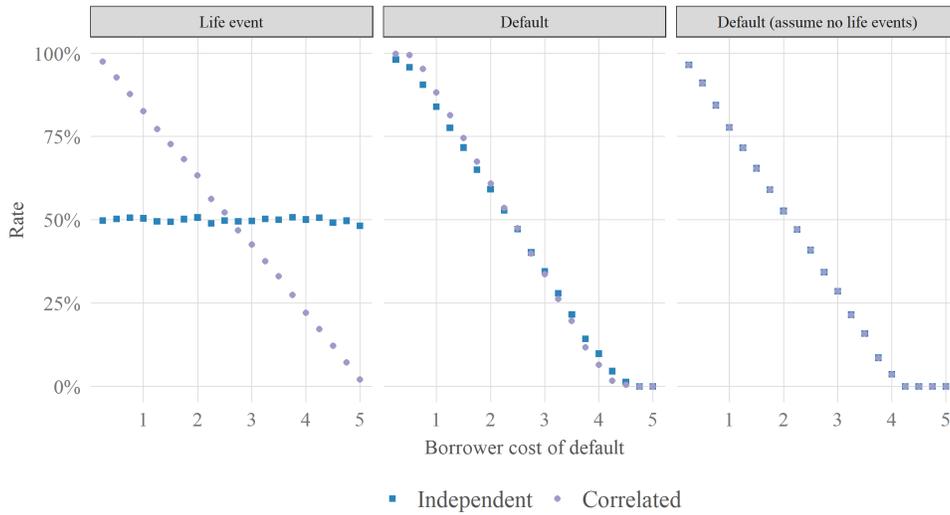
Notes: This figure reports two robustness checks on Figure 6, which uses $y - m - c_{predefault}$ as the x-variable and constructs home equity groups using an LTV cutoff of 100. Panel (a) uses an alternative x-variable $y - m - c_{subsistence}$, where $c_{subsistence}$ is a measure of the expenditure required to achieve a subsistence level of spending on non-housing consumption goods. Panel (b) uses an alternative LTV cutoff of 90, which is the cutoff used in Gerardi et al. (2018). See Section 5 for details.

Figure A-17: Campbell and Cocco (2015) Structural Model Replication



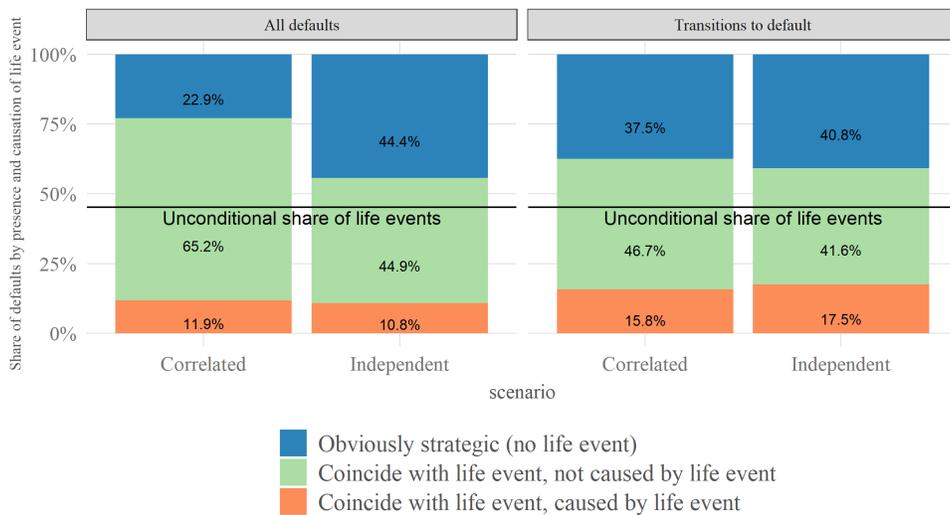
Notes: This figure shows that we can replicate the summary statistics in Table 2 of Campbell and Cocco (2015).

Figure A-18: Conditional Outcomes by Borrower Cost of Default



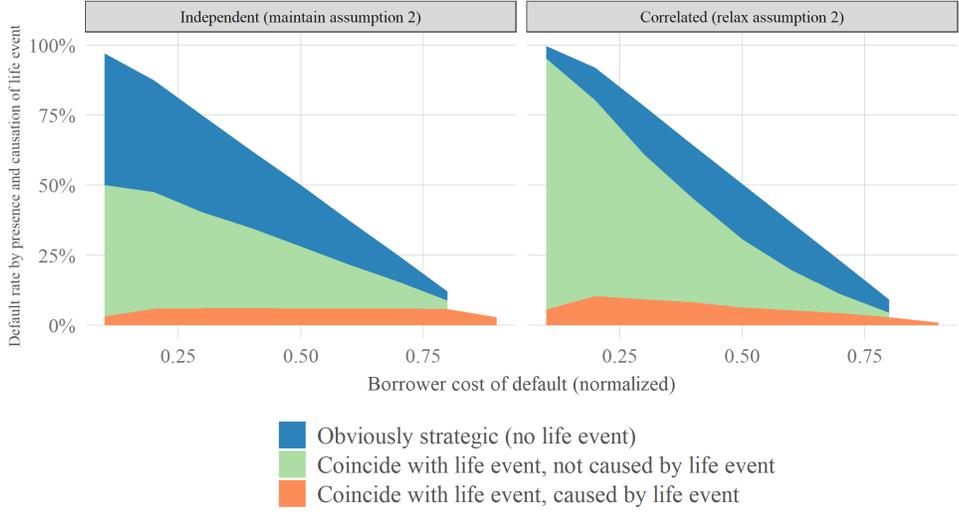
Notes: This plot shows conditional means for vintiles of the borrower cost of default (ε) under two different simulation scenarios. The left panel shows the probability of a life event, the middle panel shows the probability of default in the full simulation model, and the left panel shows the probability of default in the absence of life events. See Appendix D for details.

Figure A-19: Cause of Default by DGP and Set of Included Defaults



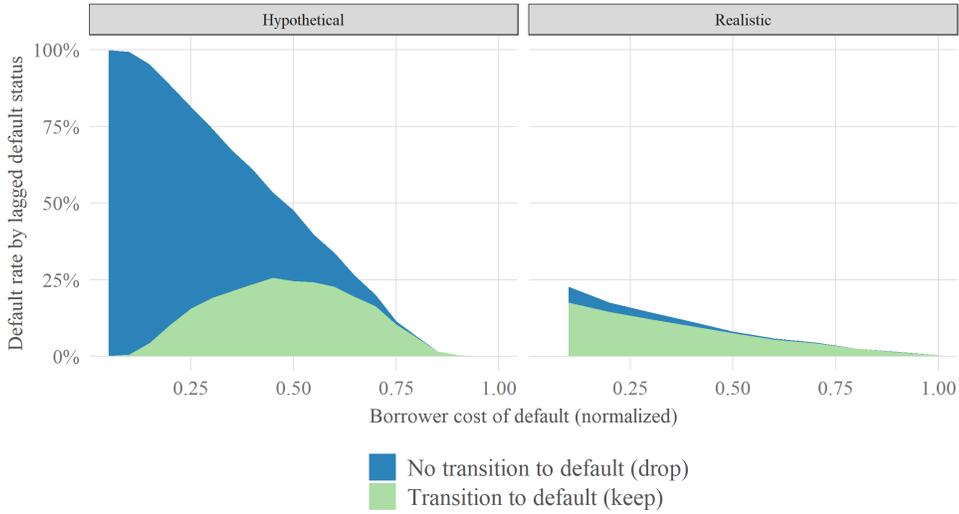
Notes: This figure shows the fraction of defaults in three groups: obviously strategic ($Y(0, 1) = 1$), coincide with life event not caused by life event ($Y(1, 1) = 1, Y(0, 1) = 1$) and coincide with & caused by life event ($Y(1, 1) = 1, Y(0, 1) = 0$) for the simulation described in Appendix D.

Figure A-20: Cause of Default by DGP and (Unobserved) Cost of Default



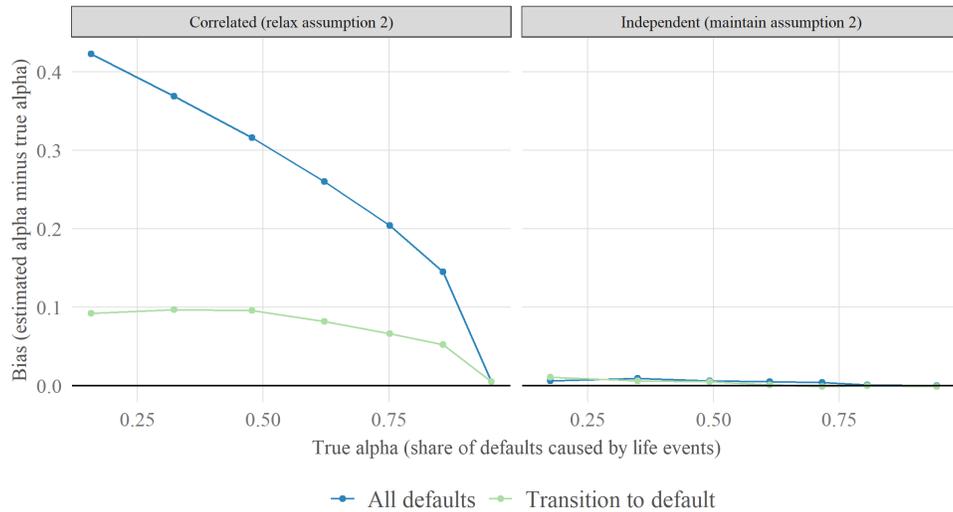
Notes: This figure disaggregates the types of default among the “all defaults” panel of Figure A-19 by the borrower cost of default ε . See Appendix D for details.

Figure A-21: Lagged Default Status by Scenario and (Unobserved) Cost of Default



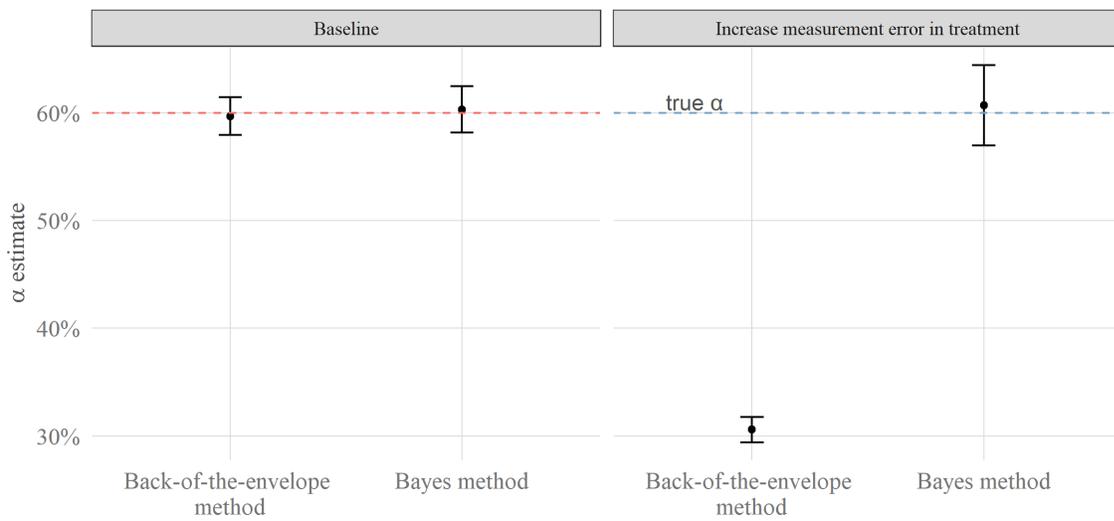
Notes: This figure disaggregates lagged default status among the “all defaults” panel of Figure A-19 by the borrower cost of default ε . See Appendix D for details.

Figure A-22: Bias in Estimated Share of Defaults Caused by Life Events (α) Scenarios



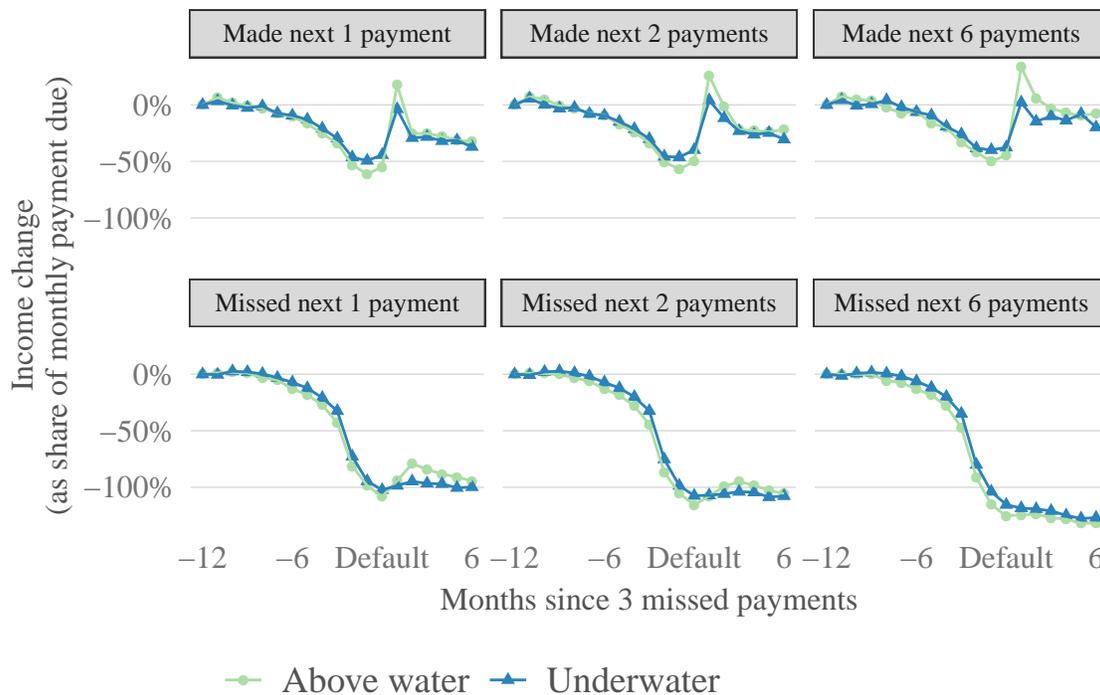
Notes: This figure shows the bias ($\hat{\alpha} - \alpha$) when studying all defaults and when narrowing the sample to just the transition to default. See Appendix D for details.

Figure A-23: Bias-Variance Trade-Off



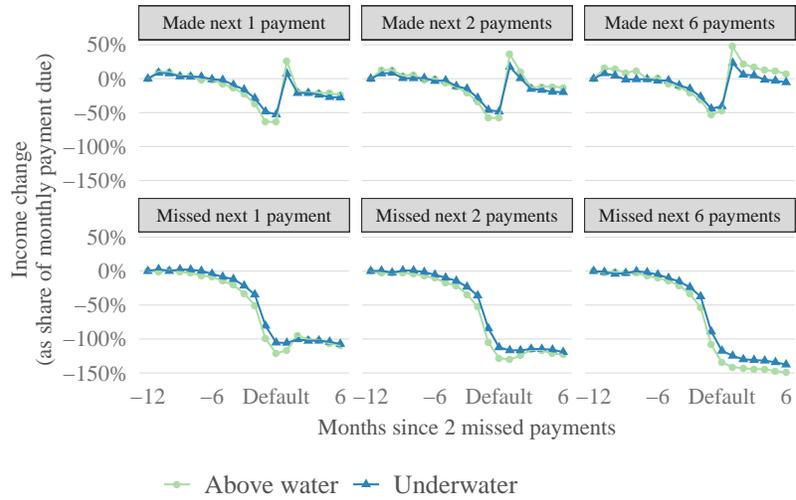
Notes: This figure shows the bias-variance trade-off between the standard “back-of-the-envelope” method and the new “Bayes” method for causal attribution. The figure reports estimates of α and a 95 percent confidence interval using both approaches within the context of a simulation. In the simulation the true α is 60 percent, denoted by the dashed horizontal line. The left panel shows the baseline case, which has minimal measurement error in treatment. In this scenario, both the “back-of-the-envelope” and “Bayes” methods are unbiased and precise. The right panel shows the noisy case, where treatment is measured with substantial error. In this scenario, attenuation bias causes the estimate from the “back-of-the-envelope” method to be biased towards zero (but still precise), whereas the “Bayes” method is unbiased (but less precise). See Appendix E for details.

Figure A-24: Income After Three Missed Payments by Subsequent Payment Behavior

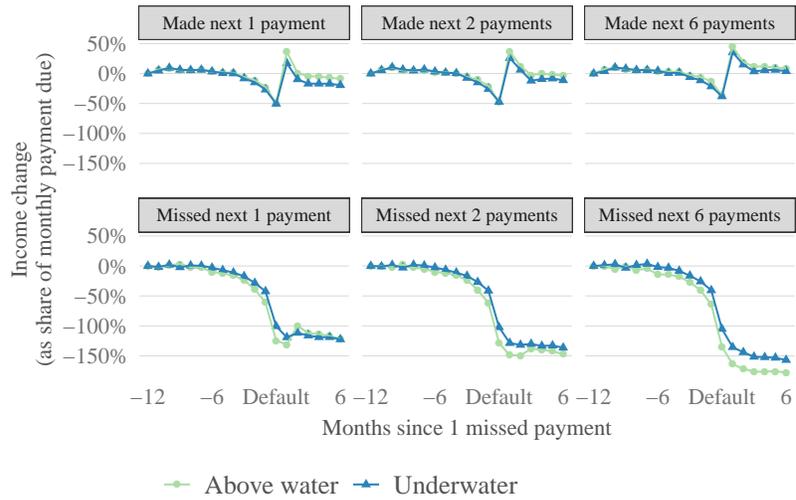


Notes: This figure analyzes the evolution of income after default (defined as three missed payments). Define $\Delta Y_t = 1$ as a deterioration in delinquency status and $\Delta Y_t = 0$ as no deterioration or an improvement. With t indexing the date of default, the top-left panel shows borrowers with $\Delta Y_{t+1} = 0$, the bottom-left panel shows borrowers with $\Delta Y_{t+1} = 1$, the top-middle panel shows borrowers with $\Delta Y_{t+1} = 0, \Delta Y_{t+2} = 0$, the bottom-middle panel shows borrowers with $\Delta Y_{t+1} = 1, \Delta Y_{t+2} = 1$, the top-right panel shows borrowers with $\Delta Y_{t+s} = 0, s \in \{1 \dots 6\}$, and the bottom-right panel shows borrowers with $\Delta Y_{t+s} = 1, s \in \{1 \dots 6\}$.

Figure A-25: Income After One or Two Missed Payments by Subsequent Payment Behavior



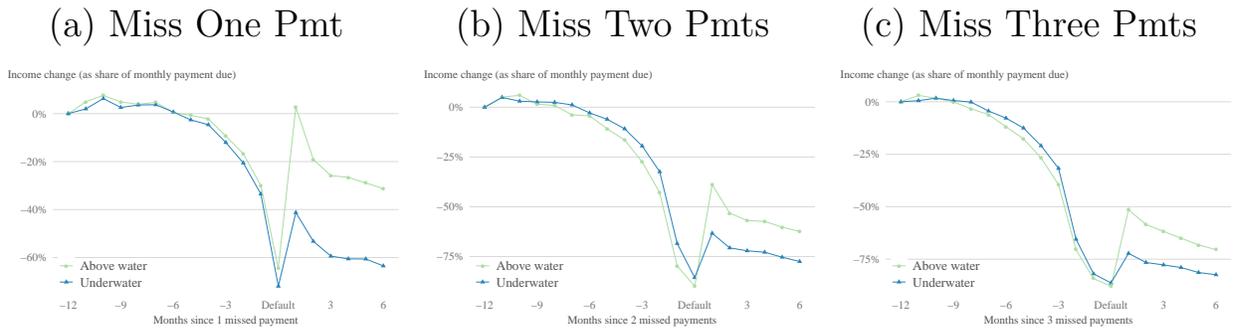
(a) After Two Payments



(b) After One Payment

Notes: This figure replicates Figure A-24 for alternative missed payments thresholds.

Figure A-26: Income After Default: Pooling Across Payment Histories



Notes: This figure shows the evolution of income both prior to and after default by alternative definitions of default, pooling across all payment paths subsequent to the default date.

Table A-1: Summary Statistics: Chase versus CRISM and McDash

Sample	Benchmark	Chase	CRISM	McDash
All mortgages	90 day delinquency rate	3.2%	3.3%	3.8%
All mortgages	Share investor	6.8%	3.9%	5.6%
All mortgages	Share primary occupant	89%	93%	91%
All mortgages	Share underwater	19%	22%	
Defaulters	Share investor	6.4%	4.3%	5.9%
Defaulters	Share primary occupant	90%	94%	92%
Defaulters	Share underwater	50%	58%	
Defaulters	Share of above water defaults with foreclosure within year	40%	49%	
Defaulters	Share of underwater defaults with foreclosure within year	45%	51%	

Notes: This table compares summary statistics regarding the matched mortgage-bank account dataset from Chase to the Credit Risk Insight Servicing McDash (CRISM) dataset in 2011. The CRISM dataset is constructed by linking credit bureau records from Equifax with mortgage servicing records from McDash. Positive and negative equity status is only observed in the linked dataset and not in McDash because it requires total mortgage debt calculated from the credit bureau data. We use 2011 as the comparison year because this is the year when U.S. house prices reached their nadir. Investor and primary occupant are reported by borrowers at mortgage origination. “Foreclosure” indicates that the mortgage servicer initiated foreclosure proceedings.

Table A-2: Income Drop From Unemployment by Home Equity

	<i>Dependent variable:</i>
	Change in income (as share of mortgage payment due)
Post UI receipt	-0.251 (0.014)
Post UI receipt * underwater	0.026 (0.026)
Observations	394,374

Notes: This table reports a regression of the income drop after unemployment by home equity. The regression is analogous to equation (7). We compare the income change in the three months after the start of unemployment (measured by the receipt of unemployment insurance (UI) benefits, as in Ganong and Noel 2019) to the income in a three-month pre-period one year before the start of unemployment (i.e. $t = \{-12, -11, -10\}$). As in equation (7), the income change is normalized by the mortgage payment due. Standard errors are clustered by mortgage.

Table A-3: Income Change for All Underwater Borrowers

<i>Dependent variable:</i>	
Change in income	
Post	-0.010 (0.003)
N mortgages	1,891,046
Observations	11,346,276

Notes: This table estimates the average income change for all underwater borrowers in months $t = \{-2, -1, 0\}$ from Figure 1. Section 4.1 provides details on how this series is constructed. The dependent variable is the ratio of monthly income to average monthly payment due in the pre-period (months $t = \{-12, -11, -10\}$). The regression specification is $\frac{Income_t}{Payment_{pre}} = \lambda + \phi \mathbf{1}(t = -2, -1, 0) + \varepsilon$. The table reports estimates for $\hat{\phi}$. Standard errors are clustered by mortgage.

Table A-4: Income Drop at Default – Straight Default and Foreclosure

	<i>Dependent variable:</i>	
	Change in income from one year before default Straight default	Foreclosure initiation
	(1)	(2)
Date of default	-1.292 (0.013)	-0.817 (0.010)
Date of default * underwater	0.159 (0.018)	0.030 (0.014)
N mortgages	68,361	96,661
Observations	410,166	579,966

Notes: This table re-estimates Column (1) of Table 4 focusing on the subsample that misses three straight mortgage payments, or using foreclosure as the alternative definition of default.

Table A-5: Share of Defaults Causally Attributable to Life Events ($\hat{\alpha}_{\text{life event}}$) Using Alternative Missed Payments Cutoffs

Months past due	$\hat{\alpha}_{\text{life event}}$ (SE)
1	1.061 (0.025)
2	0.870 (0.013)
4	0.973 (0.016)
5	1.027 (0.021)

Notes: This table reports estimates of $\hat{\alpha}_{\text{life event}}$, which is the share of defaults causally attributable to life events, in the JPMCI data using alternative months past due cutoffs. $\hat{\alpha}_{\text{life event}}$ is constructed using equation (4).

Table A-6: Share of Defaults Causally Attributable to Life Events ($\hat{\alpha}_{\text{life event}}$) Using Alternative LTV Measurement Error Correction Specifications

Category	$\hat{\alpha}_{\text{life event}}$ (SE)
2x Corelogic	0.977 (0.020)
Foreclosure penalty	0.996 (0.017)

Notes: This table reports estimates of $\hat{\alpha}_{\text{life event}}$, which is the share of defaults causally attributable to life events, in the JPMCI data using alternative LTV measurement error correction specifications. The first row assumes that the measure of the LTV error is twice as large in the Chase data as it is in the CoreLogic validation data. The second row assumes that all homes in default sell at a 9 percent discount. $\hat{\alpha}_{\text{life event}}$ is constructed using equation (4). See Section 4.3 for details.

Table A-7: Income and Assets of Defaulters by Loan-to-Value

LTV	Drop as share of income	Drop as share of mortgage payment due	Checking Balance
<100	-0.237	-0.921	1,110
100-120	-0.253	-0.860	1,201
120-140	-0.253	-0.852	1,247
140-160	-0.257	-0.895	1,301
160+	-0.261	-0.944	1,286

Notes: This table measures economic conditions at the time of default by loan-to-value (LTV) bin. The first two columns show measures of the average income drop from one year prior to default to the month of default and the third column shows mean checking account balances at the date of default. Note that this table describes borrowers at the date of default, which is different from Table 3 in the main text, which describes borrowers six months before default.

Table A-8: Distribution of Checking Account Balances of Defaulters

LTV	p10	p25	p50	p75	p90
Above water	0.5	31.1	343.6	1,289.8	3,135
Underwater	3.0	46.7	417.3	1,461.6	3,521.7

Notes: This table shows the distribution of checking account balances in dollars at the date of default for the primary analysis sample in the JPMCI data. To avoid disclosing information for any single household, the table reports pseudo-medians based on cells of at least 10 observations. Note that this table describes borrowers at the date of default, which is different from Table 3 in the main text, which describes borrowers six months before default.

Table A-9: Efficacy of Methods of Estimating α in Simulated Data

Statistic	Group/Formula	Independent	Correlated
“All periods in default” method for estimating α			
Share with life event $P(T^*)$			
	All underwater [1]	50%	50%
	Underwater: defaulters [2]	56%	77%
	Above water: defaulters [3]	100%	100%
Share of defaults caused by life events			
	$\hat{\alpha} = (2 - 1)/(3-1)$	0.11	0.54
	α	0.11	0.12
	$\hat{\alpha} - \alpha$	0.00	0.42
“Transition to default” method for estimating α			
Share with life event $P(T^*)$			
	All underwater [1]	40%	50%
	Underwater: defaulters [2]	59%	62%
	Above water: defaulters [3]	100%	100%
Share of defaults caused by life events			
	$\hat{\alpha} = (2 - 1)/(3-1)$	0.18	0.25
	α	0.18	0.16
	$\hat{\alpha} - \alpha$	0.00	0.09

Notes: This table uses simulations to explore estimates of α (the share of defaults caused by negative life events) under different assumptions about the data-generating process. In the “Independent” scenario, we assume that the probability of life events is independent of the borrower’s default cost. In the “Correlated” scenario, we assume that the probability of a life event is correlated with the borrower’s default cost. The top panel shows estimates where we study the correlates of all defaults and the bottom panel shows estimates where we study the correlates of transitions to default. See Appendix D for details.

Table A-10: Income Drop at Default by Home Equity in Actual Data

	Dependent variable: change in income	
All underwater borrowers	-0.010	0.026
Above water defaulters	-0.921	-0.295
Underwater defaulters	-0.870	-0.290
$\hat{\alpha}$	0.944	0.982
Default definition	Transition to default	All periods in default

Notes: This table shows the change in income from one year before default to the date of default as a share of the mortgage payment due in the JPMCI analysis sample. See Appendix D for details.

Table A-11: Distribution of Payment Behavior After Three Missed Payments by Home Equity

Outcome	Share of above water	Share of underwater	Difference
Made next 1 payment	0.390	0.277	0.113
Made next 2 payments	0.237	0.164	0.073
Made next 6 payments	0.094	0.068	0.026
Missed next 1 payment	0.604	0.722	-0.118
Missed next 2 payments	0.484	0.619	-0.135
Missed next 6 payments	0.266	0.382	-0.115

Note: This table shows the share of borrowers for each group in Figure A-24

Table A-12: Measures of Unemployment Among Mortgagors by Home Equity

	Unemployed	Receive UI by direct deposit
Above water	7.3%	2.0%
Underwater	12.2%	3.4%
Ratio underwater to above water	1.67	1.68
Data source	PSID	JPMCI

Note: PSID definition is head of household or spouse is unemployed at time of survey, using the 2009 and 2011 waves of the PSID. JPMCI definition is the share of borrowers who receive unemployment insurance by direct deposit in any given three month period from 2008 to 2015.

B Data appendix

B.1 JPMCI data

To be included in the analysis sample, we require that the household have an open checking account (to measure income) from one year before default through the date of default. The checking account data are available January 2007 forward. The mortgage servicing data are available through August 2015. We define the date of default as the first date a default appears in the matched checking account-servicing dataset after January 2007. To meet the requirement of having income history for the year before default, we study defaults from January 2008 through August 2015.

The unit of observation in this study is a first lien mortgage. There are 137,985 mortgages which meet this definition of default, have reliable data on payments made, have non-missing loan-to-value ratios, and have income data available for one year prior to default. There are 132,807 unique households associated with these 137,985 mortgages; this situation arises because there are a very small number of households that default on multiple first lien mortgages that are serviced by Chase.

Our primary analysis sample uses borrowers who reach delinquency of 90 days past due for the first time and have checking account data available in the 12 months prior to reaching 90 days past due. In some robustness analysis, we use alternative definitions of default: 30 days past due, 60 days past due, 120 days past due, 150 days past due, or a foreclosure initiation. For these alternative definitions of default, we rebuild the sample such that it meets the balanced panel requirement of checking account data available 12 months prior to the date of default.

After building the sample, we also take further steps to clean the data. The following variables are winsorized to the 95th percentile of positive values: monthly income, end-of-month checking account balance, monthly payment due, monthly payment made, and property value. LTV is defined as the median of non-missing LTVs between 9 and 3 months before default.

In some cases, a customer will have more than one checking account with Chase. In this case, we define income and balances summing across accounts. The crosswalk from customers to accounts is only available from 2009 forward; for income and balances prior to 2009, we use the 2009 crosswalk.

B.2 CRISM data

The CRISM data is composed of two datasets: a mortgage servicing data from McDash and credit bureau data from Equifax. The datasets are linked by the availability of a mortgage ID key in the credit data. We take a 10 percent sample of these consumers for computational reasons.

In analyzing the CRISM data, we broadly follow the data cleaning choices in Beraja et al. (2019). For each consumer, we consider their loans in both Equifax and McDash. We restrict our attention to those consumers which are first observed with less than three first liens in Equifax, and less than three of any type of second lien. Equifax only reports separate features (such as origination date, outstanding balance, origination balance) for the largest two loans, and also reports variables that reflect aggregate totals for all loans. Restricting

to those who start the panel with two or fewer loans allows us to track a third loan through changes to the total variables.

We convert the Equifax data to a loan level panel, identifying a loan by its origination date, origination amount, and lien type, and match loans in Equifax to those observed in McDash. This serves as a check on the quality of the Equifax match, and gives us more comprehensive information on second liens. We consider an Equifax loan/McDash loan pairing a match if the origination date of the Equifax loan is within one month and the origination amount is within \$10,000 of the McDash loan. If more than one McDash loan is matched, we use the date of origination, origination amount, and date of termination as tie breakers. We allow multiple Equifax loans to be matched to a single McDash loan, since occasionally loan information is revised in Equifax.

For each loan in Equifax, we consider second liens (closed-end seconds and HELOCs) that are (i) from the same consumer ID, (ii) opened along side or after after the first lien (allowing for 3 months delay in reporting), and (iii) lower in origination balance than the first lien. If a second lien is plausibly assigned to multiple first liens, it gets assigned to any first lien that has the maximum balance at origination among those first liens.

We use default information and primary balance from McDash and merge it onto second lien information as established above. This allows us to observe a measure of the cumulative unpaid balance at default. For each loan, we also observe a ZIP code and CBSA in McDash. We use this geographic information to inflate appraisals at origination by CoreLogic house price indices. We use the CBSA index where the ZIP code index is unavailable. After estimating prices and unpaid balance we have estimates of a cumulative loan-to-value at default. Additional information from the McDash occupancy field is used to classify borrowers into investors or primary occupants. Additional information from Equifax is used to observe whether borrowers are current on lines of revolving credit.

In calculating rates of foreclosure within 12 months, we consider the universe of loans that newly transition into default at a given date. We consider the share in foreclose as the proportion of these loans that are observed in foreclosure at any time from the first 90 day default to 12 months following the first 90 day default.

For the analysis in Figure A-13, we use the Equifax side of the data to measure the concurrence of default on non-HELOC revolving debt and on the primary mortgage lien. This deviates slightly from the analysis in the remaining exhibits which only uses the McDash default date. We measure the concurrence of default in a single dataset to avoid any potential issues with differences in timings between the Equifax and McDash data.

B.3 Quantifying measurement error in observed LTVs

We use two components of the CoreLogic data from 1989 to 2019: the house price indexes (to measure predicted sale values) and the deed data to measure the change in prices using 12 million home sales.

In analyzing the CoreLogic data, we broadly follow the data cleaning choices of Giacoletti (Forthcoming), with a few exceptions. First, we use a ZIP-level CoreLogic price index instead of the Zillow index that Giacoletti uses. Second, Giacoletti requires properties to also appear in the tax data; we do not, because the tax data is not required for our analysis. Third, we do not remove observations in the top and bottom 2.5 percent of errors in the main

analysis. We want to avoid understating the mass in the tails of the error distribution since this is where misclassification is most likely. Instead, we follow Kotova and Zhang (2020) in removing anomalous transactions that have aggregate appreciation or depreciation of more than 50 percent per year over their holding period. Fourth, we do not restrict our sample geographically and take any valid pair of transactions in the deed data as our sample. Finally, although we maintain Giacoletti's sample restriction regarding the earliest date of transactions to avoid issues with data quality of earlier records, we do not restrict our sample to final transactions taking place in or before 2013. This extends our sample through to the beginning of 2019.

C Proof of proposition 1

$$\begin{aligned}
\alpha_{\text{life event}} &\equiv \frac{E(Y|G = 1) - E(Y_{01}|G = 1)}{E(Y|G = 1)} \\
&= 1 - \frac{E(Y_{01}|G = 1, T^* = 0)}{E(Y|G = 1)} \\
&= 1 - \frac{P(Y = 1|T^* = 0, G = 1)}{P(Y = 1|G = 1)} \\
&= 1 - \frac{P(T^* = 0|Y = 1, G = 1)}{P(T^* = 0|G = 1)}
\end{aligned} \tag{13}$$

where the first step uses assumption 2 (random assignment of T^*), the second step uses that Y is binary, and the third step uses Bayes rule. We first analyze the numerator ($P(T^* = 0|Y = 1, G = 1)$) and then analyze the denominator ($P(T^* = 0|G = 1)$). Although neither the numerator nor the denominator are identified without further assumptions, the ratio of the two is identified using assumptions 1-4.

The law of iterated expectations implies that

$$\begin{aligned}
E(T|Y = 1, G = 1) &= P(T^* = 0|Y = 1, G = 1)E(T(0)|T^* = 0, Y = 1, G = 1) \\
&\quad + (1 - P(T^* = 0|Y = 1, G = 1))E(T(1)|T^* = 1, Y = 1, G = 1)
\end{aligned}$$

where $T(T^*, G, Y) = T(T^*)$ from assumption 3a. Re-arranging terms gives:

$$\begin{aligned}
P(T^* = 0|Y = 1, G = 1) &= \frac{E(T(1)|T^* = 1, Y = 1, G = 1) - E(T|Y = 1, G = 1)}{E(T(1)|T^* = 1, Y = 1, G = 1) - E(T(0)|T^* = 0, Y = 1, G = 1)} \\
&= \frac{E(T(1)) - E(T|Y = 1, G = 1)}{E(T(1)) - E(T(0))}
\end{aligned} \tag{14}$$

where the second equality follows from assumption 3a. This object exists because $E(T(1)) - E(T(0)) \neq 0$ by assumption 3b. We can identify $E(T(1))$ because

$$\begin{aligned}
E(T(1)) &= E(T|Y = 1, G = 0, T^* = 1)P(T^* = 1|Y = 1, G = 0) \\
&= E(T|Y = 1, G = 0)
\end{aligned} \tag{15}$$

where $P(T^* = 1|Y = 1, G = 0) = 1$ by assumption 1. Substitute equation (15) into the numerator of equation (14) to get

$$P(T^* = 0|Y = 1, G = 1) = \frac{E(T|Y = 1, G = 0) - E(T|Y = 1, G = 1)}{E(T(1)) - E(T(0))} \quad (16)$$

This expression captures the numerator of the ratio in equation (13). Applying the same logic to the denominator in the ratio of equation (13) gives

$$\begin{aligned} P(T^* = 0|G = 1) &= \frac{E(T(1)|T^* = 1, G = 1) - E(T|G = 1)}{E(T(1)|T^* = 1, G = 1) - E(T(0)|T^* = 0, G = 1)} \\ &= \frac{E(T(1)) - E(T|G = 1)}{E(T(1)) - E(T(0))} \\ &= \frac{E(T|Y = 1, G = 0) - E(T|G = 1)}{E(T(1)) - E(T(0))} \end{aligned} \quad (17)$$

where $E(T|G = 1)$ includes both underwater defaulters and non-defaulters. We take the ratio of equations (16) and (17). The denominators ($E(T(1)) - E(T(0))$) cancel, so

$$\frac{P(T^* = 0|Y = 1, G = 1)}{P(T^* = 0|G = 1)} = \frac{E(T|Y = 1, G = 0) - E(T|Y = 1, G = 1)}{E(T|Y = 1, G = 0) - E(T|G = 1)}.$$

Plugging this ratio into equation (13) gives

$$\alpha_{\text{life event}} = 1 - \frac{P(T^* = 0|Y = 1, G = 1)}{P(T^* = 0|G = 1)} = \frac{E(T|Y = 1, G = 1) - E(T|G = 1)}{E(T|Y = 1, G = 0) - E(T|G = 1)}. \blacksquare$$

Note that $E(T(0))$ cancels when computing α and so knowledge of $E(T(0))$ is not necessary for identifying α . This is why it is possible to identify the causal object α even though both the treatment effect and the probability of treatment are unknown.

D Relaxing Assumption 2

D.1 Simulate Data

We posit a statistical model of default behavior which is designed to capture the three possible theories of default: negative life events, negative equity (strategic), and double trigger. Assume a panel of borrowers, each observed for S periods. In each period, borrower default is a function of three borrower-specific variables: T_s^* is a life event, η_s is a purely temporary default shifter which governs the excess default motivations of underwater borrowers, and ε is a permanent attribute, where a high ε indicates a low cost of default. ε is intended to capture long-run attributes that affect the probability of default such as the “moral” or “stigma” cost of default. The key assumption in the simulation is that ε reflects a cost of default that is stable within the time horizon studied in the simulation. We assume that the three variables enter the default equation additively

$$Y_s = \mathbb{1}(\varepsilon + \eta_s + \beta T_s^* > a) \quad (18)$$

which allows for the possibility of interactions between the different forces (i.e., double-trigger behavior). We further assume the following distributions for the primitives:

$$\begin{aligned}\varepsilon &\sim Unif[0, \bar{\varepsilon}] \\ \eta_s &\sim Unif[0, \bar{\eta}]\end{aligned}$$

We consider two polar cases for the data-generating process for negative life events:

$$T_s^* \sim \begin{cases} Bernoulli(p) & Independent \\ Bernoulli(p \frac{2\varepsilon}{\bar{\varepsilon}}) & Correlated \end{cases}$$

Finally, relative to the framework in the paper, we treat T^* as observed and do not model T because the purpose of this simulation is to focus on causal inference regarding the role of T^* .

The goal of the simulation is to enable us to understand the impact of relaxing Assumption 2 on the paper’s estimates. In the “Independent” scenario, Assumption 2 is satisfied because life events are random. We also consider an alternative, less stringent version of Assumption 2, which is:

Assumption 2[#] (conditional exogeneity): $\{Y_s(0, 1, \varepsilon), Y_s(1, 0, \varepsilon), Y_s(1, 1, \varepsilon)\} \perp T_s^* | G, \varepsilon$

In the “Correlated” scenario, Assumption 2 is not satisfied because there is a correlation between the individual cost of default ε and the probability of a life event, but Assumption 2[#] is satisfied.

The prior literature provides no obvious guidance on the sign of the omitted variable bias (whether ε is positively or negatively correlated with the probability of a life event T_s^*).² We focus in our simulations on the negative correlation case because this paper’s primary conclusion is that almost no defaults are strategic and that conclusion might change when Assumption 2 is relaxed to allow for a negative correlation (i.e., a negative correlation is what would lead us to understate the share of strategic default). We allow for a severe form of omitted variable bias: the probability of a life event ranges from 100 percent for the borrowers with the lowest cost of default to nearly 0 percent for the borrowers with the highest cost of default. We normalize the probability of a life event by $2/\bar{\varepsilon}$ so that the overall probability of a life event is similar under the two assumptions about the data-generating process for T^* . Furthermore, we adopt parameters such that the bias in our estimate is maximized, which happens when the share of defaults causally attributable to life events

²One strand of papers posits that strategic default is most common among borrowers who are investors and have low attachment to the home in question, and therefore may have low private costs of default (Albanesi, De Giorgi and Nosal, 2017). These borrowers, who previously had enough resources to buy properties beyond their primary residence, may be less likely to have life events than the average borrower. If that was true, then we might expect a *positive* correlation between default costs and probability of a life event. A second strand posits that strategic default is most common among subprime borrowers who have low attachment to credit markets overall, and therefore may also have low private costs of default since their credit score is already low (Mayer et al., 2014). These borrowers may be more likely to have life events than the average borrower if low credit scores are more common for borrowers with high latent unemployment risk (perhaps arising from past spells of unemployment). If this was true, then we might expect a *negative* correlation between default costs and the probability of a life event.

is small. We show that, even with these two extreme assumptions, the overall bias in our estimator is small. Furthermore, when we adopt more realistic assumptions, the bias is fully mitigated.

D.2 Analyze Simulated Data

D.2.1 Analysis of Using All Periods in Default to Estimate α

When we use our method to analyze every period in which a borrower is in default, our methodology does well when Assumption 2 holds but does poorly in some parameterizations when Assumption 2 is relaxed. In this subsection, we deviate from the paper by defining $Y_s = 1$ as any period with a default. We provide this example to illustrate how omitted variable bias can lead to misleading conclusions about a causal relationship, but note that this is *not* the method we use in the paper. We discuss bias from the actual method in the next section.

In the case where Assumption 2 holds ($Y_s(0, 1) \perp T_s^* | G$), the probability of a life event T_s^* is independent of the cost of default ε .³ We simulate a parameterization of the model where $\bar{\varepsilon} = 5, \bar{\eta} = 4, p = 0.5, \beta = 0.5, a = 5$ with 10,000 borrowers each of whom live for 10 periods. A sample size of 10,000 borrowers for the simulation is sufficiently large that we can ignore issues of sampling variation and focus on the key question of causal inference. We show robustness to alternative parameterizations in Section D.2.2. We assume that the probability of a life event is independent of the other model parameters ($T_s^* \sim \text{Bernoulli}(p)$) and call this the “Independent” scenario.

Figure A-18 shows in the first panel that, consistent with this assumption, the average probability of a life event does not vary with the cost of default in the Independent scenario. It also shows in the second panel that the probability of default is highest for those with the lowest cost of default (those with high ε) and in the third panel that the default rate would change very little if there were no life events. We estimate equation 3 (substituting T^* for T). Table A-9 shows that we estimate $\hat{\alpha}_{xsec} = 0.11$, which means that 11 percent of defaults are caused by life events. This matches the true parameter value in the simulation ($1 - E(Y_s(0, 1) | Y_s = 1, G = 1)$). This is to be expected, because we proved that this set of conditional expectations identifies the causal object of interest in Proposition 1.

Next, we consider a case where we relax Assumption 2 by introducing substantial omitted variable bias and show that estimator applied to all periods with a default does poorly. Specifically, we assume that the probability of a life event is higher for people with a low stigma cost of default ($T_s^* \sim \text{Bernoulli}(p \frac{2\varepsilon}{\eta})$) and call this the “Correlated” scenario. Figure A-18 shows that in this scenario, the probability of a life event varies from 100 percent for people with the lowest cost of default to 0 percent for people with the highest cost of default. By design, the extent of omitted variable bias is extreme in this situation: the correlation of ε and T_s^* is greater than 0.5. In this case, we estimate in Table A-9 that $\hat{\alpha}_{xsec} = 0.54$, when in fact the true α is 0.12.

The estimator does poorly because of the standard intuition of how omitted variable bias can lead an analyst to overstate the strength of a causal relationship. The data in this simulation feature borrower life events and borrower default that coincide, but life events

³To be precise, Assumption 2 is that $Y(0, 1) \perp T^* | G$, while here we add a subscript s to capture the panel dimension of the simulation. We otherwise ignore the panel dimension of the data in this section.

have little causal impact on default. To clarify how this affects our estimates of α , Figure A-19’s left panel shows that the “Correlated” scenario shows excess mass in the green bar relative to the “Independent” scenario. This means that there are a substantial share of defaults that have a life event but are not caused by a life event. A higher rate of life events among defaulters than among the general population leads to an (incorrect) inference that life events are causing many defaults.

D.2.2 Analysis of Transitions to Default to Estimate α

We redo the analysis from Section D.2.1, limiting attention to the subsample where $Y_{t-1} = 0$. This is the methodology we use in Section 4 in the paper. In the language of the model in the paper, which does not have a time dimension, $Y = 1$ is the *transition* to default. This is the natural outcome of interest in our setting. Every loan is current at the time of origination, and so every default requires a transition from not defaulting to defaulting. Understanding borrower default therefore requires understanding why borrowers transition into default. This is why the prior literature on mortgage default often evaluates default as time-to-failure through the lens of a hazard model.

This estimand exactly identifies the parameter of interest when Assumption 2 holds and shows substantial improvement in the scenario where Assumption 2 is relaxed. We analyze the same simulated data as in the previous section. In the “Independent” scenario, Table A-9 shows that the simulation estimate matches the true parameter value. Again, this is to be expected because of Proposition 1. In the “Correlated” scenario, the magnitude of the bias is substantially reduced; it is 1/5th as large as when we analyze all defaults.

Focusing on the transition to default mitigates the potential bias from relaxing Assumption 2 in our simulation because it drops borrower-periods where permanent heterogeneity causes misleading conclusions. These are periods after the initial transition to default for the borrowers with a low cost of default (regardless of whether they have experienced a life event) and a high probability of negative life events. Figure A-20’s left panel shows that the defaults that coincide with a life event but would have happened without it are concentrated among borrowers with a low permanent cost of default. Figure A-21 shows that narrowing the sample to just examine transitions to default drops most of the misleading default-periods. Although the extent of the bias could differ under alternative data-generating processes, we note that we have deliberately selected a data-generating process where omitted variable bias is severe (because the correlation between the cost of default ε and the probability of a life event T^* is quite strong).

Fixed effects models offer a helpful, if inexact, analogy for why analyzing transitions to default leads to minimal bias even in the presence of unobserved heterogeneity. Had Y_s been linear in the latent confounder ε , even after we relax Assumption 2, a first-difference specification would have enabled us to exactly identify α .⁴ Because Y_s is binary, the identification results from the continuous case do not hold. However, the intuition that focusing on transitions differences out the latent type carries over to this setting.

⁴Suppose we modify equation 18 such that $Y_s = \varepsilon + \eta_s + \beta T_s^*$. Note that in the “Correlated” simulation scenario above, a conditional version of exogeneity holds ($Y_s(0, 1, \varepsilon) \perp T_s^* | G, \varepsilon$). First-differencing gives $E(Y_s - Y_{s-1} | \varepsilon, \eta_1 \dots \eta_T) = \beta(T_s^* - T_{s-1}^*)$. In a model with homogeneous treatment effects, β gives the average causal effect of T_s^* on Y ; together with the probability of treatment $P(T_s^*)$, this is sufficient to identify α .

We conduct additional simulations to verify the robustness of our conclusions about bias across the support of the parameter of interest. Specifically, relative to the base simulations described above, we progressively reduces the importance of permanent heterogeneity (by decreasing $\bar{\psi}$ to 2.5) while simultaneously increasing the effect of live events on default (increasing β to 3).⁵ These changes drive α up to about 0.80. Finally, to raise α even further, we lower the probability of life events p to 0.1 and lower $\bar{\eta}$ to 2.75, which raises α to 0.94.

The results are shown in Figure A-22 and contain three lessons. First, defining $Y = 1$ as the transition to default greatly reduces (but does not eliminate) bias at most values of α . Second, the estimates are similar from using all periods with a default and using transitions to default when α is close to 1. To understand why the estimates are similar, see the panel labeled “Realistic” in Figure A-21, which shows that most defaults were *not* preceded by a default in the previous period. The two samples (all defaults and transitions to default) therefore study similar samples and yield similar conclusions. Third, there is no evidence of bias in the simulation when α is close to 1.

The results in Figure A-22 suggest a further empirical test using actual data. The figure shows that the estimates using all defaults and using transitions to default are similar when either of two conditions is satisfied: (i) when Assumption 2 is satisfied or (ii) when α is close to 1. We therefore re-implement our empirical methodology using every month in which a borrower is in default, instead of just using the first month and show the results in Table A-10. The estimate of $\hat{\alpha}$ is 0.982 when we study all periods where a borrower is in default as compared to 0.944 when we study only the transition to default. Interpreted through the lens of the simulations, the similarity of the two estimates suggests that the bias from relaxing Assumption 2 is limited.

Based on both the simulations which relax Assumption 2 and the alternative empirical estimate that uses all defaults, it appears that our conclusions about the prevalence of strategic default are robust to relaxing Assumption 2.

E Bias-Variance Trade-off

This section uses a simple simulation to explore the bias-variance trade-off for the new method of causal attribution proposed in Section 2.2 relative to the standard method described in Section 2.1. We are interested in measuring the fraction of underwater mortgage default causally attributable to life event treatment T^* . Assume that all borrowers have a 50 percent probability of binary treatment, i.e. $T^* \sim \text{Bernoulli}(0.5)$. Assume that 25 percent of underwater borrowers (those with $G = 1$) receive a binary strategic default shifter S , so that

$$S = \begin{cases} 0 & \text{if } G = 0 \\ \text{Bernoulli}(0.25) & \text{if } G = 1 \end{cases} \quad (19)$$

⁵Our identification method requires the presence of some abovewater defaults. In the simulation, this requires that $\bar{\psi} + \beta > a$. Had we decreased $\bar{\psi}$ with no change to another parameter, then this condition would not be fulfilled and no abovewater borrower would ever default. We therefore include offsetting increases to β so as to ensure that some abovewater default persists in the simulation, as it does in the data.

Borrowers default if they receive a negative life event, a strategic default shock, or both, i.e.

$$Y = \mathbb{1}(T^* + S \geq 1).$$

Given these parameters, 100 percent of underwater borrowers with a negative life event default, 25 percent of underwater borrowers without a negative life event default (due to the strategic motive), and 62.5 percent of all underwater borrowers default. Because only 25 percent of all underwater borrowers would default in the absence of life event treatment, the true $\alpha_{\text{life event}}$ for underwater borrowers from equation (1) is

$$\alpha_{\text{life event}} \equiv \frac{0.625 - 0.25}{0.625} = 0.6.$$

In other words, life events are a necessary condition for 60 percent of underwater default in this setup.

Next, we introduce a measure of treatment that is noisy. In practice, we often do not observe the treatment itself (e.g., the life event T^*) but rather a noisy proxy (e.g., the change in income, which we denote by T). In the simulation, we assume that the observed income change is centered at zero, falling on average for those with negative life events and rising on average for those without, but is observed with a normally distributed measurement error $\varepsilon_{\sigma^2} \sim N(0, \sigma^2)$, which gives:

$$T = 0.5 - T^* + \varepsilon_{\sigma^2}.$$

We consider two measurement error scenarios. In the baseline case, we assume only a small amount of measurement error with a standard deviation of 0.05, i.e. $\varepsilon_{0.0025} \sim N(0, 0.0025)$. In the noisy scenario, we assume a higher degree of measurement error, i.e. $\varepsilon_{0.25} \sim N(0, 0.25)$. We simulate each parameterization of the model with 10,000 above water borrowers and 10,000 underwater borrowers and assume the econometrician can observe (Y, T, G) for all borrowers in the simulation, but not T^* . In each case, we can compare the standard “back-of-the-envelope” method for estimating $\alpha_{\text{life event}}$ from Section 2.1 (which we denote here as $\tilde{\alpha}_{\text{life event}}$) to the new “Bayes” method proposed in Section 2.2 (which we denote here as $\hat{\alpha}_{\text{life event}}$). The results are summarized in Figure A-23.

E.1 Baseline scenario

The standard “back-of-the-envelope” method to estimating $\alpha_{\text{life event}}$ from equation (2) requires three inputs: an estimate of the average treatment effect on default, the probability of treatment, and the probability of default. Because treatment is randomly assigned, to obtain the average treatment effect the researcher would simply regress default on the noisy measure of treatment, i.e.

$$Y = \lambda + \beta T \tag{20}$$

and use $-\hat{\beta}$ as the average treatment effect (since treatment leads to an observed income *decrease*). In the baseline simulation, we calculate $-\hat{\beta} = 0.743$ for underwater borrowers, close to the true causal impact of treatment for underwater borrowers of 0.75. There is minimal attenuation bias in this estimate because the noise in T is minimal. Furthermore, the estimate of the causal impact is precise, with a standard error of 0.006. The researcher infers the probability of treatment from the share with an observed income decline, which

is 50.5 percent ($P(\hat{T}^*) = 0.505$), and observes an underwater default rate $E(\hat{Y}) = 0.628$. Plugging these in to equation (2) gives

$$\tilde{\alpha}_{\text{life event}}^{\text{baseline}} = \frac{(0.743)(0.505)}{0.628} = 0.597,$$

which, as with the causal impact, is close to the true $\alpha_{\text{life event}}$ because measurement error is minimal. We calculate a standard error using the delta method of 1 percent.

We can also implement the new ‘‘Bayes’’ method for estimating $\alpha_{\text{life event}}$ from Section 2.2. Implementing this method using equation (3) also requires three inputs: the average observed income change for underwater defaulters ($E(T|Y = 1, G = 1)$), the average observed income change for abovewater defaulters ($E(T|Y = 1, G = 0)$), and the average observed income change for all underwater borrowers ($E(T|G = 1)$). In the simulation, the average change in income for underwater defaulters is -0.304, the average change for abovewater defaulters is -0.501 (all these borrowers experience a negative life event, which leads on average to an observed income decrease of 0.5), and the average change of income for all underwater borrowers is -0.005 (since just over half receive a negative life event which leads to an average income change of -0.5 and just under half receive no negative life event, which leads to an average income change of 0.5). Plugging these values into equation (3) gives

$$\hat{\alpha}_{\text{life event}}^{\text{baseline}} = \frac{(-0.304) - (-0.005)}{(-0.501) - (-0.005)} = 0.603,$$

which is also nearly identical to the true $\alpha_{\text{life event}}$. We calculate a standard error using the delta method of 1.1 percent.

The point estimates and confidence intervals for both $\tilde{\alpha}_{\text{life event}}^{\text{baseline}}$ and $\hat{\alpha}_{\text{life event}}^{\text{baseline}}$ are depicted visually in the left panel of Figure A-23.

E.2 Noisy scenario

We implement the same approaches for measuring $\tilde{\alpha}_{\text{life event}}$ and $\hat{\alpha}_{\text{life event}}$ in the noisy scenario. Because T is measured with significant noise, there is now substantial attenuation bias in $\hat{\beta}$ from equation (20). We estimate $-\hat{\beta} = 0.381$, substantially below the true causal impact of 0.75. However, despite the bias in the estimate, because the outcome in equation (20) is measured just as precisely as in the baseline scenario, the *precision* of the estimate is just as good: the standard error on the estimate is again 0.006. The increased bias but similar precision translate to an estimate of $\tilde{\alpha}_{\text{life event}}^{\text{noisy}} = 0.301$, about half as large as the true $\alpha_{\text{life event}}$.

This trade-off is reversed with the Bayes method. The Bayes method estimates the income changes, which are now much noisier. However, they are unbiased. Thus, we find $\hat{\alpha}_{\text{life event}}^{\text{noisy}} = 0.607$, close to the true $\alpha_{\text{life event}}$, but measured much less precisely because the variance of each measured income change is now much larger. The standard error for the $\hat{\alpha}_{\text{life event}}^{\text{noisy}}$ nearly doubles, to 1.9 percent.

The point estimates and confidence intervals for both $\tilde{\alpha}_{\text{life event}}^{\text{noisy}}$ and $\hat{\alpha}_{\text{life event}}^{\text{noisy}}$ are depicted visually in the right panel of Figure A-23. Attenuation bias from measurement error in treatment causes the estimate from the ‘‘back-of-the-envelope’’ approach to be biased towards zero (but still precise), whereas the new ‘‘Bayes’’ approach is unbiased (but less precise).

F Income after default

This appendix investigates the evolution of income *after* default. We find that payment and income dynamics after default are closely linked. Defaulters who resume making payments also have an income increase back to pre-default levels, while defaulters who do not resume making payments (and therefore fall behind further) continue to have depressed incomes.

We define additional notation for this analysis. Let t index dates and Y_t index delinquency. Let $Y_t \in \{0, 30, 60, 90, 120, 150, 180, 210, 240, 270, \text{Foreclose}\}$ where the numeric values denote the number of days that the loan is past due and if the lender begins foreclosure proceedings, then the number of days is censored and instead Y_t is recorded as “Foreclose.” Finally, define

$$\Delta Y_t \equiv \begin{cases} 1 & Y_t > Y_{t-1} \text{ or } Y_t = \text{Foreclose}, Y_{t-1} \neq \text{Foreclose} \\ 0 & Y_t \leq Y_{t-1} \end{cases}$$

Following the logic of our main empirical design (which conditions on payment behavior), we begin by conditioning on payment behavior for the first payment due after default. Figure A-24 shows the evolution of income for borrowers after default, using our baseline definition of three missed payments ($Y_t = 90, Y_{t'} \leq 60 \quad \forall t' < t$). The date of default is therefore t . In the top left panel we analyze borrowers whose delinquency does not deteriorate further at time $t + 1$ ($Y_{t+1} \leq 90$) and therefore made at least one month’s worth of payments at date $t + 1$. We refer to this group as having $\Delta Y_{t+1} = 0$. Income for the remaining borrowers, who miss the next payment due ($\Delta Y_{t+1} = 1$), is shown in the bottom-left panel. These two groups partition the full set of possible payment outcomes in the month after default. After default, we find that borrowers who make a payment also experience an average income recovery, while the average income of borrowers who miss the next payment remains depressed. These patterns do not differ by home equity.

Similar patterns appear as we study additional payment history beyond the first payment due after default: borrowers who continue to make payments experience an income recovery while borrowers who continue to miss payments have income that remains depressed. The middle column of Figure A-24 shows borrowers who made the next 2 payments ($\Delta Y_{t+1} = 0, \Delta Y_{t+2} = 0$) or missed the next two payments ($\Delta Y_{t+1} = 1, \Delta Y_{t+2} = 1$). The right column shows borrowers who made the next 6 payments ($\Delta Y_{t+s} = 0, s \in \{1 \dots 6\}$) or missed the next 6 payments ($\Delta Y_{t+s} = 1, s \in \{1 \dots 6\}$). Unlike the left column, these columns with “made/missed next x payments” are no longer a full partition of payment outcomes, but nevertheless document income in scenarios of particular interest.

Two lessons emerge from Figure A-24. First, the basic “neutrality” result—that the income drop is similar for above water and underwater defaulters—holds when we look at specific payment histories after default. Second, income changes track payment histories, providing further evidence of the tight link between income and default.

Both the neutrality result and also the conclusion that income changes track payment histories extend to other definitions of default. Figure A-25 repeats the same analysis as Figure A-24, but instead defines default as two missed payments ($Y_t = 60, Y_{t'} \leq 30 \quad \forall t' < t$) and one missed payment ($Y_t = 30, Y_{t'} = 0 \quad \forall t' < t$).

Against the backdrop of the neutrality results described above, we note that there *are* two well-known empirical patterns of differences between underwater and above water borrowers. The first pattern is that underwater borrowers are unconditionally more likely to experience negative life events in any period. This is the central finding in Bhutta, Dokko, and Shan.⁶ This finding also holds using one common measure of life events (unemployment) in both datasets analyzed in this paper. We find that receipt of unemployment is almost twice as common for underwater than for above water borrowers, both in the PSID and in Chase.

The second pattern is that the rate of transition from less severe delinquency to more severe delinquency (“the roll rate”) is higher for underwater borrowers ($P(\Delta Y_1 = 1|G = 1) > P(\Delta Y_1 = 1|G = 0)$). This pattern is apparent in prior empirical research (see for example Mayer et al. 2014), and also holds in our sample. Table A-11 shows the shares for each group in Figure A-24. After 90-day default, 72 percent of underwater borrowers miss the next payment while 60 percent of above water borrowers miss the next payment.

Another interesting way to describe the data is to document the path of average income after default, pooling across subsequent payment histories. Given the neutrality result shown in Figure A-24 ($E(T_{PD}|\Delta Y_1 = 1, G = 1) \approx E(T_{PD}|\Delta Y_1 = 1, G = 0)$ and $E(T_{PD}|\Delta Y_1 = 0, G = 1) \approx E(T_{PD}|\Delta Y_1 = 0, G = 0)$ where T_{PD} is the change in income “post default” relative to 12 months prior to default) and the difference in roll rates by home equity, we expect that income after default will be more depressed for underwater borrowers than for above water borrowers. We expect that this will hold because of the law of iterated expectations

$$E(T_{PD}|G) = \underbrace{E(T_{PD}|\Delta Y_1 = 1, G)}_{\text{very negative, similar for } G \in \{0,1\}} + \underbrace{P(\Delta Y_1 = 1|G)}_{\text{higher for } G=1} \underbrace{E(T_{PD}|\Delta Y_1 = 0, G)}_{\text{less negative, similar for } G \in \{0,1\}} + \underbrace{P(\Delta Y_1 = 0|G)}_{\text{lower for } G=1}$$

which implies that $E(T_{PD}|G = 1) < E(T_{PD}|G = 0)$. Figure A-26 confirms that, as expected, the pooled income after default is lower for underwater borrowers.

Figures A-24 and A-26 suggest that the higher rate of negative life events for underwater borrowers (pattern 1 above) could be the cause of the higher roll rates for underwater borrowers (pattern 2 above). Figure A-26 shows a difference by home equity in income after default (not conditioning on payment history), while Figure A-24 shows no difference in income after default (when conditioning on payment history). If controlling for differences in payment histories eliminates differences in observed cash flows by home equity, a natural explanation is that the differences in cash flows were causing the differences in payment histories.

⁶This is also consistent with evidence in Bernstein (2019) and Gopalan et al. (Forthcoming), who find that borrowers with negative equity are more likely to suffer income declines because of constrained mobility and financial distress.

G Stigma cost of default

We are interested in finding the change in average per-period consumption (Δc) that would cause a change in utility equal to a one-time stigma cost:

$$u(\bar{c} + \Delta c) = u(\bar{c}) + \textit{stigma}.$$

We know that utility from consumption takes the following form:

$$u(c_t) = \sum_{t=1}^T \beta^{t-1} \frac{c_t^{1-\gamma}}{1-\gamma}$$

Assuming constant consumption (\bar{c}) across all periods, and letting $\gamma = 2$, as in Campbell and Cocco, we get

$$u(\bar{c}) = \frac{1}{\bar{c}} \sum_{t=1}^T -\beta^{t-1} = \frac{1}{\bar{c}} \lambda$$

where $\lambda = -\frac{1-\beta^T}{1-\beta}$. To do so we invert the utility function to get consumption as a function of utility

$$\bar{c}(u) = \frac{\lambda}{u}$$

We then use this formula to find the percent change in \bar{c} that would come from a change in average utility equal to stigma. In doing so, λ cancels out and we are left with the percent change in consumption as a function of average expected lifetime utility and stigma:

$$\begin{aligned} \%change &= \frac{\bar{c}(u + \textit{stigma}) - \bar{c}(u)}{\bar{c}(u)} \\ &= \frac{\frac{\lambda}{u + \textit{stigma}} - \frac{\lambda}{u}}{\frac{\lambda}{u}} \\ &= \frac{-\textit{stigma}}{u + \textit{stigma}} \end{aligned}$$

Plugging in the average expected lifetime utility in the period the mortgage is originated (-2.63), this formula replicates the Campbell and Cocco result ($\textit{stigma} = -0.05 \implies \%change = -1.9\%$) and gives $\%change = -25.5\%$ at our estimated value of $\textit{stigma} = -0.90$ that best fits the JPMCI data.

We therefore calculate that a borrower whose behavior matches the Campbell and Cocco model estimated to fit the JPMCI data would give up about \$100,000 in consumption to avoid default. Ignoring bequests, non-housing consumption is equal to income minus housing expenditures. Non-housing consumption in the Campbell and Cocco model is income minus housing expenditures. Average income is \$43,300 in Table II panel (d) and housing expenditures are 40 percent of income at the time of default, so non-housing consumption is \$26,000 per year. The consumption equivalent the borrower would give up is therefore \$6,600 per year. Discounting at 3 percent and summing over 20 years, this gives a present value of \$100,000.