WHY DO BORROWERS DEFAULT ON MORTGAGES?*

Peter Ganong and Pascal Noel

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 compare these theories in part because negative life events are measured with error. We address this measurement error using a comparison group of borrowers with no strategic-default motive. Our central finding is that only 6% of underwater defaults are caused exclusively by negative equity, an order of magnitude lower than previously thought. We then analyze the remaining defaults. We find that 70% are driven *solely* by negative life events (i.e., cash flow defaults), while 24% are driven by the *interaction* between negative life events and negative equity (i.e., double-trigger defaults). Together, the results provide a full decomposition of the theories underlying borrower default and suggest that negative life events play a central role. *JEL Codes:* G21, G51, G41.

*We are particularly grateful for many fruitful conversations with Kanav Bhagat that helped shape the analysis in this article. We thank John Campbell, João Cocco, and Marco Giacoletti for generously sharing code and for very helpful comments. We further thank João Cocco, Angus Foulis, Amir Kermani, and Paul Willen for serving as discussants on this paper. We also thank Neil Bhutta, Gabriel Chodorow-Reich, Adam Guren, Kyle Herkenhoff, Peter Hull, Erik Hurst, Koichiro Ito, Anil Kashyap, Ben Keys, David Matsa, Neale Mahoney, Atif Mian, Jack Mountjoy, Mikkel Plagborg-Møller, 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 for helpful comments. We thank Ari Anisfeld, Rei Bertoldi, Therese Bonomo, Guillermo Carranza Jordan, Timotej Cejka, Lei Ma, Roshan Mahanth, Michael Meyer, Liam Purkey, Peter Robertson, John Spence, Nicolas Wuthenow, and Katie Zhang 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 deidentified data assets that are selectively available 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 article, 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, the Kathryn and Grant Swick Faculty Research Fund, and the Fujimori/Mou Faculty Research Fund at the University of Chicago Booth School of Business.

The Quarterly Journal of Economics (2023), 1001–1065. https://doi.org/10.1093/qje/qjac040.
Advance Access publication on October 13, 2022.

I. 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." (Bernanke 2008)

Mortgage defaults soared during the Great Recession, precipitating the worst financial crisis since the Great Depression. As Ben Bernanke explained, a key challenge facing lenders and policy makers 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.

The challenge of distinguishing between negative equity and negative life events during the Great Recession was the latest chapter in a long-standing 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 article is to help distinguish between these theories. Our central finding is that only 6% of underwater defaults are caused exclusively by negative equity, an order of magnitude lower than previously thought. We then further decompose the remaining defaults. We find that 70% are driven *solely* by negative life events (i.e., cash flow defaults), while the remaining 24% 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 review article (Foote and Willen 2018). The question has remained central partly 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 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 data set with details about a household's financial circumstances at the time of default. In a review article published over two decades later, Foote and Willen (2018) call for constructing almost exactly the same data set.

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%)? 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 an alternative empirical method. First, to overcome the data limitation,

1. A 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 IV.C.

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 data set has 3.2 million borrowers, which is much 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] 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 compare the evolution of income for underwater borrowers to this benchmark.

We find that 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 long-standing theories of default in terms of potential outcomes

2. It may initially be surprising that any borrowers with positive equity ever default (which we define in 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 (DeFusco and Mondragon 2020; Boar, Gorea, and Midrigan 2022). As a result, missed mortgage payments are ubiquitous for borrowers with positive equity (Low 2018).

TABLE I
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

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

(Rubin 1974). We summarize the model types in Table I and describe them formally in Section II. We assume that either negative equity or a negative life event is necessary for default. We label a default as (i) "strategic" when negative equity is a necessary and sufficient condition for the default; (ii) "cash flow" when a negative life event is necessary and sufficient; and finally, (iii) "double-trigger" when both negative equity and a negative life event are necessary.[3] The model provides a tractable framework for decomposing the share of defaults attributable to these three theories.

Second, we show that the evolution of income before default can be used to identify the share of strategic defaults in 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

3. The label of "strategic default" as a default 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 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 V.

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 the above-water defaulters.

We combine these two ingredients using a procedure similar to reverse regression. The standard approach to causal inference puts the outcome on the left side and treatment on the right side of a regression equation. However, as is well known, this approach suffers from attenuation bias when treatment is measured with error. Naively regressing default on a noisy measure of life events will lead researchers to underestimate their importance in driving defaults. We therefore use reverse regression 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. We condition on default and put the change in bank account income—which is our measure of life events—on the left side.

Our estimating equation has a simple interpretation. 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 nondefaulters, then we would classify all their defaults as strategic.

When we interpret our reduced-form empirical findings through the lens of this econometric approach, our estimates show that at most 6% of underwater defaults are strategic. In other words, we find that negative life events are a necessary condition for 94% of mortgage defaults, so 94% of defaults must be either cash flow or double-trigger. Although concern about borrowers walking away from their homes solely because of negative equity was widespread (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 0% to 5%.

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%. Indeed, our test detects little evidence of strategic default until borrowers have LTVs above 200%. Fourth, our baseline specification follows much of 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 I) from the double-trigger defaults (row 3 in Table I). This requires one additional ingredient: the causal effect of negative equity on default. We estimate this causal effect by adopting methods from the literature based on long-run regional variation in house price cyclicality (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] Indeed, we do find evidence for substantial double-trigger behavior: both triggers are necessary for 24% of underwater defaults. However, we find that pure cash flow defaults, driven *solely* by negative life events, are also important: they account for 70% 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 I 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% of defaults (strategic plus double-trigger) while eliminating life events would prevent 94% of defaults (cash flow plus double-trigger).[5]

4. See 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.

5. These estimates complement a literature that analyzes the effect of variation in penalties for debt nonrepayment on bankruptcy (Dávila 2020; Indarte 2020; Gross et al. 2021).

These results help interpret the strong cross-sectional correlation between negative equity and default. It is well documented that default rates are higher for borrowers with negative equity than for borrowers with positive equity. Is this indicative of strategic default? Our results suggest that it is not. Instead, negative equity borrowers are both substantially more likely to experience cash flow shocks and also more likely to default conditional on experiencing such a shock.

Our finding of the near absence of strategic default contrasts with prior estimates that between 30% and 70% of Great Recession defaults were strategic (Guiso, Sapienza, and Zingales 2013; Bhutta, Dokko, and Shan 2017; Gerardi et al. 2018). To help explain *why* our estimates are so much lower, we reanalyze publicly available survey data from the Panel Study of Income Dynamics (PSID). Although the administrative bank data are key for obtaining precise estimates, we show that the survey data also yield 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 also contrasts with existing structural models that predict substantial strategic default by deeply underwater borrowers. We use a benchmark

6. The 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 Defaults." Yet as researchers have acquired more detailed data and 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.

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 of the model can reconcile them. 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% decrease in the constant-equivalent consumption stream. This could reflect financial costs through reduced access to credit, or nonfinancial costs such as a borrower's attachment to their current home (i.e., their idiosyncratic private valuation of the home may be greater than the market's valuation). We discuss potential sources of this cost in more detail in Section VI. 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 macroprudential regulation, the origins of the 2008 financial crisis, bankruptcy and foreclosure policy, and optimal mortgage security design.[8] These types of models must

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

8. For example, see Corbae and Quintin (2015); Mitman (2016); Garriga and Hedlund (2020); Kaplan, Mitman, and Violante (2020); Campbell, Clara, and Cocco (2021); Greenwald, Landvoigt, and Van Nieuwerburgh (2021); Guren, Krishnamurthy, and Mcquade (2021); Chodorow-Reich, Guren, and McQuade (2022); and Diamond and Landvoigt (2022).

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 defaults. 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 article proceeds as follows. Section II describes our econometric framework. Section III describes the two data sets with income, home equity, and default (administrative bank data and PSID survey data). Section IV 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 V replicates our finding of little strategic default using the PSID. Section VI explores implications for modeling mortgage default. Section VII concludes.

II. Econometric Framework

II.A. A Potential-Outcomes Model of Mortgage Default

1. *Environment and Assumptions.* We assume that there exists a population distribution (T^*, G, Y, T) where T^* is the treatment of a negative life event, G is the group status of having negative equity, Y is the transition to mortgage default, and T is the change in bank account income relative to one year ago, which is a candidate noisy measure of T^*. The first three variables are binary. Assume that there exists a potential-outcome function $Y(T^*, G)$, so that each individual has four potential outcomes: $Y(0, 0), Y(1, 0), Y(0, 1)$, and $Y(1, 1)$. In addition, 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 a negative life event or negative equity): $Y(0, 0) = 0$ with probability one.

This assumption says that the outcome of mortgage default requires either negative equity or a negative 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 of thinking of this assumption is that, following the terminology from Bernanke's speech enumerating a number of different types of life events, we define a "life event" ($T^* = 1$) as anything that causes an above-water borrower to default. Our central causal object of interest, which we define below, 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 nonnegative equity causes of default.[9]

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).[10] 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. Yet in practice, above-water default is quite common because there are substantial frictions to accessing home equity for borrowers in financial distress (DeFusco and Mondragon 2020; Boar, Gorea, and Midrigan 2022). Distressed borrowers may alternatively attempt to sell their home, but there are frictions in this process as well (Guren 2018; Gilbukh and Goldsmith-Pinkham 2021). In essence, these frictions make home equity less liquid in the short term than it is in the long term. In Online Appendix C.1, we discuss these frictions in more detail as well as evidence on the prevalence of and reasons for above-water default and foreclosure.

9. One 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% increase in new defaults. This means that, in the quarter that the program was implemented, 10% of defaults for borrowers missing consecutive payments were motivated by a desire to get a mortgage modification.

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

TABLE II

POTENTIAL-OUTCOMES MODEL OF THE MORTGAGE DEFAULT DECISION

			$Y(\underbrace{T^*}_{\text{life event}},$	$\underbrace{G}_{\text{negative equity}})$	
Type	Decision rule	$Y(0, 0)$	$Y(1, 0)$	$Y(0, 1)$	$Y(1, 1)$
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 II. The fact that there are three potential-outcome types arises from Assumptions 1 and 2 and our focus on defaulters who are underwater. See Section II for details.

ASSUMPTION 2 (Monotonicity): $Y(1, 1) \geqslant Y(1, 0)$, $Y(1, 1) \geqslant Y(0, 1)$.

This assumption says that negative life events and negative equity each make a borrower weakly more likely to default.

2. *Economic Interpretation.* This framework allows us to separate default behavior in terms of three potential-outcome types corresponding to the three theories for default in the prior literature. We summarize these types in Table II.

i. First, there are "strategic" 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 ($Y(0, 1) = Y(1, 1) = 1$) but would not default solely because of a life event ($Y(1, 0) = 0$).
ii. Second, there are "cash flow" 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 ($Y(1, 0) = Y(1, 1) = 1$) but would not default solely due to negative equity ($Y(0, 1) = 0$).
iii. Third, there are "double-trigger" 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 ($Y(1, 1) = 1$), but would not default if only one trigger was present ($Y(0, 1) = Y(1, 0) = 0$).

Under Assumptions 1 and 2, all underwater defaulters are one of these three potential-outcome types.[11] Let ST, CF, and DT denote the share of underwater defaulters with the strategic, cash flow, and double-trigger potential-outcome types, respectively. Because all underwater defaulters have one of these three types,

$$ST + CF + DT = 1. \tag{1}$$

We now discuss how we identify the share of each of these types.

II.B. Causal-Attribution Estimands

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

$$\alpha \equiv \frac{E[Y] - E[Y(0)]}{E[Y]}, \tag{2}$$

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 central goal of our analysis is to estimate the fraction of underwater defaults that are causally attributable to negative life events. This can be captured by measuring the fraction of underwater defaults that would be eliminated in the absence of

11. The fact that there are three potential-outcome types of interest arises from Assumptions 1 and 2 and our focus on defaulters who are underwater. See Online Appendix C.2 for details.

life events ($T^* = 0$). Define

$$(3) \qquad \alpha_{\text{life event}} \equiv \frac{E[Y(T^*, 1)|G = 1] - E[Y(0, 1)|G = 1]}{E[Y(T^*, 1)|G = 1]}$$
$$= 1 - ST$$
$$= CF + DT,$$

where the second line follows because strategic defaulters are the only type with $Y(0, 1) = 1$, and thus the only defaulters who would remain if life events were eliminated, and the third line follows from equation (1). Thus, $\alpha_{\text{life event}}$ captures both the combined share of cash flow and double-trigger defaults (for which life events are a necessary condition) and also, by process of elimination, the share of strategic defaults (for which life events are not a necessary condition).

A secondary goal of our analysis is to look within the defaults for which life events are a necessary condition and distinguish between the cash flow and double-trigger defaults. This can be accomplished by asking the same type of counterfactual question as in our primary analysis, but for negative equity instead of negative life events. Specifically, we ask: what fraction of underwater defaults would be eliminated in the absence of negative equity? This share can be denoted as:

$$(4) \qquad \alpha_{\text{negative equity}} \equiv \frac{E[Y(T^*, 1)|G = 1] - E[Y(T^*, 0)|G = 1]}{E[Y(T^*, 1)|G = 1]}$$
$$= 1 - CF$$
$$= ST + DT.$$

Without negative equity, there would be no strategic or double-trigger defaults (for which negative equity is a necessary condition), but there would still be cash flow defaults. Thus, $\alpha_{\text{negative equity}}$ captures the combined share of strategic and double-trigger defaults and also, by process of elimination, the share of cash flow defaults. Finally, we can recover the double-trigger share by substituting equations (3) and (4) into equation (1): $DT = \alpha_{\text{life event}} + \alpha_{\text{negative equity}} - 1$.

II.C. Causal-Attribution Identification

1. *Standard Identification Approach to Causal Attribution.* The standard empirical method for estimating α parameters is sometimes called a back-of-the-envelope calculation. To illustrate this approach, for expositional simplicity we focus just on underwater borrowers and assume a single treatment T^* with potential-outcome function $Y(T^*)$. Because $Y = Y(1)T^* + Y(0)(1 - T^*)$, α can be rewritten as proportional to the product of the average treatment effect and the probability of treatment:

$$\alpha = \frac{E[Y] - E[Y(0)]}{E[Y]} = \frac{E[Y(1)T^*] + E[Y(0)(1 - T^*)] - E[Y(0)]}{E[Y]}$$

$$= \frac{E[Y(1) - Y(0)|T^* = 1]P(T^*)}{E[Y]}. \quad (5)$$

Researchers typically estimate a treatment effect, multiply by an estimate of the probability of treatment $(P(\hat{T}^*))$ and divide by an estimate of the average level of the outcome $(E(\hat{Y}))$.

Applied to analyzing the fraction of mortgage default attributed to life events $\alpha_{\text{life event}}$, 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 (5) to hypothetically conclude that 50% of default is attributable to life events. However, in settings where treatment is particularly difficult to observe, measurement error can bias estimates of α because of attenuation bias in the estimate of treatment effects and because the probability of treatment may be unknown.

The measurement error challenge makes it difficult to estimate $\alpha_{\text{life event}}$ using the standard approach. 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 can't we use 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 effect 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 $P(T^*)$. Gerardi et al. (2015) estimates that among mortgagors (defaulters and nondefaulters), the probability of a negative life event ranges from 4% under a stringent definition of a large decline in income to 57% under a broad definition that includes several types of shocks.

The bulk of this article is devoted to estimating $\alpha_{\text{life event}}$ using a method that is unbiased even in the presence of measurement error. In contrast, 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. In Section IV.E we report estimates of $\alpha_{\text{negative equity}}$.

2. *The Reverse-Regression Approach to Causal Attribution.* We estimate $\alpha_{\text{life event}}$ using two ingredients. The first ingredient is a group whose outcome $Y = 1$ must have been caused by treatment T^*: above-water defaulters, whose defaults we assume must be caused by a negative life event (Assumption 1). The second ingredient is a noisy measure of treatment T, which we use to substitute for unobserved true treatment T^*. We combine these ingredients using reverse regression. This requires two additional assumptions common in the measurement error literature.

ASSUMPTION 3 (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 home equity. This assumption is standard in the measurement error literature such as in the classical errors-in-variables (CEV) framework in Wooldridge (2010) and in the literature studying the motivations for default (Bhutta, Dokko, and Shan 2017; 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 (Hembre 2018; Laufer 2018; Schelkle 2018; Kaplan, Mitman, and Violante 2020). In such models, Assumption 3 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.[12]

12. 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 Online Appendix C.5, we relax Assumption 3 by allowing for correlated, latent heterogeneity in the probability of a life event and the probability of

This assumption allows for two 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 the findings in Gerardi et al. (2018) that underwater borrowers are more sensitive to income shocks than above-water borrowers ($E[Y(1, 1) - Y(0, 1)] > E[Y(1, 0) - Y(0, 0)]$).

ASSUMPTION 4 (Noisy measure of treatment): (i) $T(T^*, G, Y) = T(T^*)$ and $\{T(0), T(1)\} \perp (T^*, Y, G)$, and (ii) $E(T(1)) \neq E(T(0))$

Assumption 4i says that the potential-outcome function for T is orthogonal to the other variables in the model: life event T^*, home equity G, and default 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).

Assumption 4ii 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 4 implies that when a life event does occur, above- and underwater borrowers have the same average decline in income. This allows us to use the income of above-water defaulters, who always have $T^* = 1$ by Assumption 1, to learn about $P(T^*)$ for underwater defaulters. Online Appendix C.3 discusses testable implications of Assumption 4 and shows that the data are consistent with these assumptions.

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.

13. This is also consistent with evidence in Bernstein (2021) and Gopalan et al. (2021), which find that borrowers with negative equity are more likely to suffer income declines because of constrained mobility and financial distress.

PROPOSITION 1: Under the environment described above and Assumptions 1, 2, 3, and 4,

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

$$= \frac{E[T|Y=1, G=1] - E[T|G=1]}{E[T|Y=1, G=0] - E[T|G=1]}. \tag{6}$$

Proof: See Online Appendix C.4.

The standard approach to causal attribution puts the outcome Y on the left side and treatment T^* on the right 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^*).

The key step in the proof relies on reverse regression to correct for attenuation bias. The classic example of reverse regression is a researcher who seeks to measure whether an employer is discriminating against their female employees in setting wages. The ideal test is to regress wages on employee productivity T^* 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 T) are typically observed. Reverse regression puts the credentials T on the left side, wages on the right 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.

As with the classic reverse-regression application, we correct for attenuation bias by moving the noisy variable T from the right side to the left side of the regression equation. Reverse regression also requires that what is usually the left-side variable (wages in the classic case, mortgage default in our case) be measured without error (Goldberger 1984). With administrative data on mortgage default, this condition should be satisfied in our application. Unlike the classic reverse-regression case, we also require that 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 VI. Finally, we note that the benefits and costs of reverse regression can be framed in terms of a bias-variance trade-off, which we explore using a simulation in Online Appendix C.6.

3. *Causal Attribution of Mortgage Defaults.* The formula is easy to interpret because it relies on comparing the change in income for underwater defaulters (our group of interest) to two benchmarks. To see this comparison clearly and ease interpretation as we take the formula to the data, we can annotate equation (6) as:

$$\alpha_{\text{life event}} = \frac{\overbrace{E(T\,|Y=1,G=1)}^{\text{group of interest: underwater defaulters}} \qquad -E(T\,|G=1)}{\underbrace{E(T\,|Y=1,G=0)}_{\alpha=1\text{ benchmark: above-water defaulters (life events)}} - \underbrace{E(T\,|G=1)}_{\alpha=0\text{ benchmark: all underwater (strategic)}}}. \tag{7}$$

The first benchmark describes what it would look like in the data 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, that is, $E(T|Y = 1, G = 0)$. 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 describes what would it look like in the data if defaults were driven exclusively by negative equity. 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 nondefaulters. Hence the $\alpha_{\text{life event}} = 0$ benchmark is the change in income for all underwater borrowers, including defaulters and nondefaulters, that is, $E(T|G = 1)$. This benchmark arises intuitively from a 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" (Foster and Van Order 1984, 353).

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 nondefaulters. Thus, if we observe the same income drop for underwater defaulters and nondefaulters (i.e., $E(T|Y = 1, G = 1) = E(T|G = 1)$), then our assumptions imply that negative life events play no role in the default decision ($\alpha_{\text{life event}} = 0$).

Finally, if the income change for underwater defaulters is in between the two benchmarks, then the share of defaults causally attributed to life events is between zero and one.

III. Data

III.A. Linked Income-Servicing Data

Our primary analysis uses a novel administrative data set from Chase that links checking-account records and mortgage-servicing records. These records are linked and then deidentified by Chase.[14] 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 posttax 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, although 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. Observed income provides a useful summary measure of many life events that would be difficult to observe directly in the data, such as an injury or other health shock that limits a borrower's ability to work.

14. See Farrell et al. (2017) and Farrell, Bhagat, and Zhao (2018) for JPMCI research using this linked data set.

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.[15] Estimated home value is constructed using the standard procedure of inflating the 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 a 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 IV.C.

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

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 data set 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 86,693 above-water defaults and 52,519 underwater defaults from 2008 through 2015 (see Online Appendix B for details on sample construction). We also analyze the evolution of income for the universe of underwater borrowers, both defaulters and nondefaulters (657,053 borrowers). Finally, for our analysis of the impact of negative equity on default for computational reasons we subset to a 15% random sample of the full linked data set (above-water and underwater, defaulters and nondefaulters). This subsample includes 451,590 borrowers and 13.5 million borrower-months.

We assess the representativeness of the Chase data by comparing the characteristics of our sample to those from other

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

TABLE III

Distribution of Home Equity and Default versus Benchmarks

	Default rate (%)			Share (%)		
LTV bin	JPMCI	CRISM	PSID	JPMCI	CRISM	PSID
LTV $>$ 100	8.7	8.7	7.8	19.2	21.9	10.6
80 $<$ LTV $\leqslant$ 100	3.2	3.0	2.5	23.9	26.8	26.4
LTV $\leqslant$ 80	1.2	1.1	1.1	56.9	51.4	63.0

Notes. This table compares the distribution of home equity and default rates for mortgage borrowers in the JPMorgan Chase Institute (JPMCI) data to the Credit Risk Insight Servicing McDash (CRISM) data set, and the Panel Study of Income Dynamics (PSID) in 2011. We use 2011 as the comparison year because this is the year when U.S. house prices reached their nadir. Online Appendix Table A-11 reports the same statistics for 2009. Default is defined as three missed mortgage payments.

data sets commonly used in the mortgage default literature. Table III shows that the gradient of default rates with respect to home equity (the key cross-sectional relationship that motivated economists to become concerned about potentially strategic behavior) is similar to the gradient in Credit Risk Insight Servicing McDash (CRISM) and the Panel Study of Income Dynamics (PSID). Moreover, the distribution of home equity is also similar.[16] We discuss these data sets in more detail below. Online Appendix Table A-1 shows that the Chase data are also similar to McDash, CRISM, and MBA in terms of origination and performance characteristics. Finally, Online Appendix Table A-2 shows that Chase borrowers are slightly older than those in the PSID.

III.B. Other Data Sets

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 a parallel analysis using the best available public-use data on income and mortgage default, which is survey data from the PSID. The PSID records pretax income y and consumption c in the calendar year before the survey. A supplement administered in 2009, 2011, and 2013 records housing costs (mortgage, property taxes, and insurance) m, home equity, and default at the time of the survey. We follow the sampling choices from Gerardi et al. (2018): we drop households with LTV $>$ 250%, and we require that the head of household is in the labor force and between the ages of 24 and

16. Gerardi et al. (2018) note that negative equity is relatively underreported in the PSID.

65. We discuss the strengths and weaknesses of the PSID data compared to the JPMCI data in Section V.

For our analysis on the share of defaults causally attributable to negative equity, we conduct a parallel analysis using CRISM data. McDash is a data set of origination and servicing records that covers about 70% 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% of outstanding mortgage balances during this time period. Our CRISM analysis on the causal effect of negative equity uses a 1% random sample of borrowers with first-lien mortgages who appear in the data between 2008 and 2015. This sample includes 388,000 borrowers (see Online Appendix B.2 for details on sample construction).

IV. 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 II, 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 effect of negative equity on default to distinguish between the remaining cash flow and double-trigger defaults.

IV.A. Research Design for Strategic Defaults

We compare underwater defaulters' income to two benchmarks using equation (7). In this section we discuss each of the two benchmarks. In Section IV.B we use all three objects to estimate the share of defaults causally attributable to life events.

The green triangles in Figure I show the evolution of income for above-water defaulters in the 12 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 from an event study by focusing specifically on income data *prior* to default. This choice is motivated by this article'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

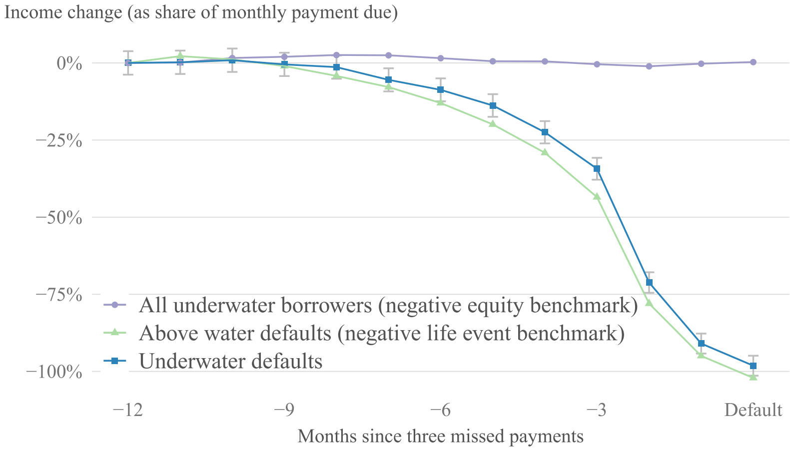

FIGURE I

Monthly Evolution of Income in the Year Prior to Mortgage Default

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 defaulters in comparison to two benchmarks: income for all underwater borrowers in circles, which captures the negative equity (strategic) benchmark, and income for above-water defaulters in triangles, which captures the negative life event benchmark. The gray error bars show the 95% 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 IV for details.

in monthly income relative to the average mortgage payment due one year before default.[17] 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. Online Appendix Figure A-2 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.[18]

17. This normalization facilitates the interpretation of the point estimates in terms of number of mortgage payments due and the comparison of our data to the model in Section VI. Online Appendix Figure A-1 shows similar patterns (indeed, even less evidence of strategic default) when normalizing by prior income rather than by payment due.

18. However, 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

Figure I also shows the evolution of income for all underwater borrowers (both defaulters and nondefaulters). This series provides the benchmark for what income would look like if all defaults were strategic. We construct this series by reweighting average income by month to match the realized distribution of default dates across all months in our sample. 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 $\{s_{default} - 12, s_{default} - 11, \ldots, s_{default}\}$. 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:

$$(8) \qquad Income^{AllUnderwater} = \sum_s Income_s^{AllUnderwater} w_s.$$

To capture average income of all underwater borrowers 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 $t \in \{-12, -11, \ldots, 0\}$. Figure I 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 IV shows that underwater borrowers have slightly higher income levels, bank account balances, and mortgage payments due than above-water borrowers. Instead, Assumption 4 requires that income declines by the same amount conditional on a life event. Online Appendix Table A-12 shows that this does indeed hold for one life event we can reliably observe: unemployment.

IV.B. 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 I. The figure shows that income falls for underwater defaulters nearly as much as for above-water

event in the context of our model. Furthermore, payment changes directly affect borrowers' ability to pay.

TABLE IV
SUMMARY STATISTICS BY HOME EQUITY

Variable	Above water	Underwater
Combined loan-to-value ratio (%)	75	125
Bank account monthly income ($)	4,284	4,543
Bank account balance ($)	1,529	1,769
Property value ($)	248,478	222,072
Monthly mortgage payment due ($)	1,141	1,363
Age	49	47
Share with joint deposit account	0.4	0.45
N	86,693	52,519

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 the 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.

defaulters. The gray error bars show the 95% 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 "preperiod" well before default ($t = \{-12, -11, -10\}$) and a three-month period around the time of default ($t = \{-2, -1, 0\}$). We regress

$$\frac{Income_t}{\bar{Payment}_{pre}} = \lambda + \kappa \mathbf{1}(LTV > 100) + \underbrace{\gamma \mathbf{1}(t = -2, -1, 0)}_{\text{above-water drop at default}}$$

$$+ \underbrace{\beta \mathbf{1}(t = -2, -1, 0) \times \mathbf{1}(LTV > 100)}_{\text{difference for underwater}} + \varepsilon, \tag{9}$$

where $\bar{Payment}_{pre}$ is the average payment in the preperiod, computed separately for above- and underwater defaulters. Table V, column (1) shows that the above-water income drop $\hat{\gamma}$ is −0.928. This means that monthly income falls on average by 93% as a share of the monthly payment in the month of mortgage default and the two months prior. Table V also shows that $\hat{\beta}$ is 0.057.

Applying the framework from Section II to our regression estimates, we find that 94% of underwater defaults are causally attributable to life events. Equation (7) requires three inputs to

TABLE V

INCOME DROP AT DEFAULT BY HOME EQUITY

	Dependent variable:				
	Change in income from one year before default				
	Mean (1)	Median (2)	p25 (3)	p75 (4)	Mean (w/ME correction) (5)
Date of default	−0.928	−0.764	−0.572	−1.267	−0.925
	(0.008)	(0.009)	(0.010)	(0.016)	(0.010)
Date of default × underwater	0.057	−0.035	−0.057	0.063	0.045
	(0.012)	(0.015)	(0.016)	(0.022)	(0.018)
N mortgages	139,212	139,212	139,212	139,212	139,212
Observations	835,272	835,272	835,272	835,272	835,272

Notes. This table describes the income drop at default in the JPMCI data. The sample uses a balanced panel with a preperiod (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 preperiod. The regression specification is $\frac{Income_t}{\bar{Payment}_{pre}} = \lambda + \kappa(LTV > 100) + \gamma \mathbf{1}(t = -2, -1, 0) + \beta \mathbf{1}(t = -2, -1, 0) \times LTV > 100 + \varepsilon$, as shown in equation (9). 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 change in the 50th, 25th, and 75th percentiles of the income distribution, respectively. Column (5) repeats column (1), but using the measurement error (ME) correction for LTV from equation (12). 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.

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 Online Appendix Table A-3). We estimate that

(10)
$$\hat{\alpha}_{\text{life event}} = \frac{(\hat{\gamma} + \hat{\beta}) - \hat{\varphi}}{\hat{\gamma} - \hat{\varphi}} = \frac{-0.928 + 0.057 - (-0.009)}{-0.928 - (-0.009)} = 94\%$$

with a standard error, computed using the delta method, of 1%, as shown in Table VI, Panel A. The 95% confidence interval on the share of underwater defaults causally attributable to negative life events ($\hat{\alpha}_{\text{life event}}$) ranges from 91% to 96%.

We use the potential-outcomes model in Table II 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% of underwater defaults are strategic.

TABLE VI

Share of Defaults Causally Attributable to Life Events ($\hat{\alpha}_{\text{LIFE EVENT}}$)

Category	$\hat{\alpha}_{\text{life event}}$ (std. err.)
Panel A: Baseline estimate and robustness for all borrowers	
Baseline	0.938 (0.013)
Change from $t = -12$ to $t = 0$	0.961 (0.016)
Foreclosure	0.961 (0.017)
LTV measurement error correction (all sales)	0.951 (0.016)
LTV measurement error correction (foreclosures)	1.002 (0.019)
Panel B: Subsample heterogeneity	
Year	
2008	0.914 (0.058)
2009	0.948 (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	0.932 (0.079)
LTV	
101–120	0.927 (0.015)
121–140	0.918 (0.023)
141–160	0.963 (0.034)
161–180	1.067 (0.052)
181–200	1.036 (0.087)
201–220	0.887 (0.089)
221+	0.865 (0.098)
Mortgage type	
Fixed	0.931 (0.015)
Adjustable	1.016 (0.029)
Nonrecourse states	1.021 (0.027)
Three consecutive missed payments	0.871 (0.013)
Subprime borrowers	0.905 (0.042)

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 (7). Panel A shows estimates for all borrowers. The first row is the baseline specification, using the regression in equation (9) and defining default as three missed payments. The second row implements equation (9) 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 and fifth rows implement the two-sample IV measurement error correction procedure for LTV described in Section IV.C using the distribution of house price errors from all sales and from foreclosure sales, respectively. Panel B shows estimates for specific subsets of borrowers using the baseline specification. See Sections IV.B and IV.C for details.

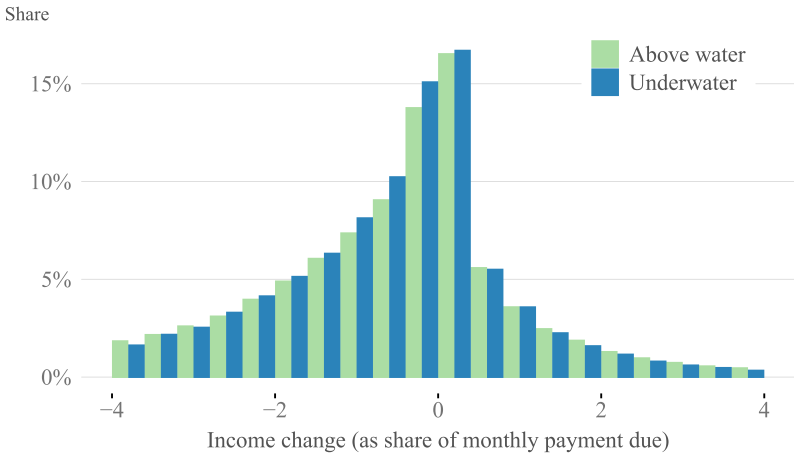

FIGURE II

Distribution of Income Change in the Year Prior to Mortgage Default

This figure shows the full distribution of the income change in Figure I from the three months around default relative to the first three months of the series (as in equation (9)) and is truncated at −4 and 4 to improve readability.

IV.C. Robustness for Strategic Defaults

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

1. *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 (4%) if we focus instead on the income drop just in the month of default, as shown in Table VI, Panel A.

2. *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 II shows that the entire distribution of the change in monthly income is similar for above- and underwater borrowers.[19] Quantitative estimates in Table V columns (2)–(4) similarly show that the change in the 25th, 50th, and 75th percentiles of the income distribution are similar for above- and underwater borrowers. Finally, event study patterns for changes in different percentiles of the

19. Online Appendix Figure A-3 shows the same lesson using the cumulative distribution function.

income distribution are also similar, as shown in Online Appendix Figure A-4.

The distribution of income changes is also useful for evaluating a testable implication of Assumption 4. 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 that 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. The 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 II shows no evidence of such heterogeneity.

One notable feature of Figure II 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.[20]

3. *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

20. The 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.

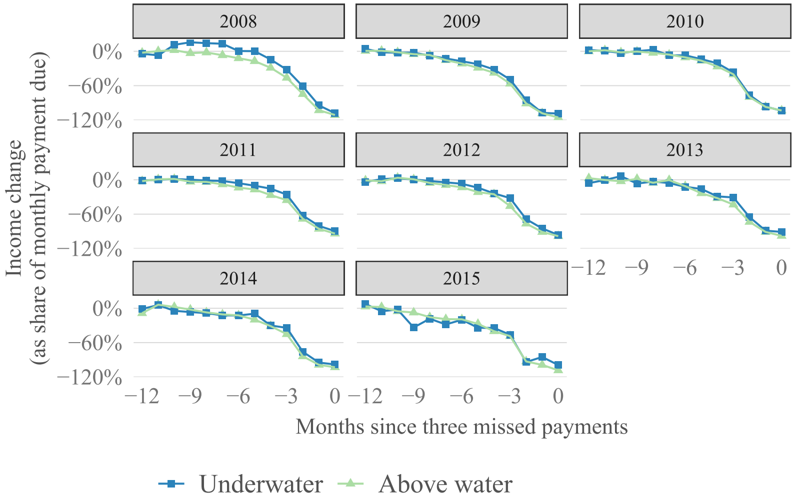

FIGURE III

Heterogeneity in Income Drop before Default by Year

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 IV.C for details.

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 III shows that income declines are similar for above- and underwater borrowers in every year from 2008 through 2015. Likewise, Table VI, Panel B shows that the estimates of $\alpha_{\text{life event}}$ are similar across different years.

4. *Mortgage Type*. We find similar patterns across different mortgage types. Online Appendix Figure A-5 plots income before default by home equity separately for borrowers with fixed rate and adjustable-rate mortgages, showing similar patterns. Table VI, Panel B shows estimates of $\alpha_{\text{life event}}$ for each mortgage type.

5. *Nonrecourse States*. Our results are also not sensitive to whether the default occurs in a recourse or a nonrecourse state.

Online Appendix Figure A-6 and Table VI, Panel B show that there is no evidence of more strategic default in states with nonrecourse mortgage debt, where defaulting may be more financially advantageous. However, this may not be surprising for two reasons. First, Ghent and Kudlyak (2011) report that deficiency judgments against borrowers in recourse states are exceedingly rare. Second, Guiso, Sapienza, and Zingales (2013) report no difference between borrowers in recourse and nonrecourse states in their subjective estimates of the probability that lenders would pursue them after a default.

6. *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. Online Appendix Figure A-7 shows that the income patterns before foreclosure are similar for above- and underwater borrowers. Table VI, Panel A shows that our estimate for the share of strategic default using this definition of default (4%) is also similar to our baseline estimate.[21] Furthermore, although our baseline specification defines default as three missed payments, Online Appendix Figure A-8 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. Online Appendix Table A-5 reports estimates of $\alpha_{\text{life event}}$ for each of these missed payment specifications. In addition, in Online Appendix D.1 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.

7. *Heterogeneity by LTV Ratio*. Our results are also not sensitive to the choice of LTV threshold. Figure IV 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 (LTV $<$ 60) to those with substantial negative equity (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 VI, Panel B shows estimates of $\alpha_{\text{life event}}$

21. Online Appendix Table A-4 replicates the regression in equation (9) using foreclosure as the definition of default.

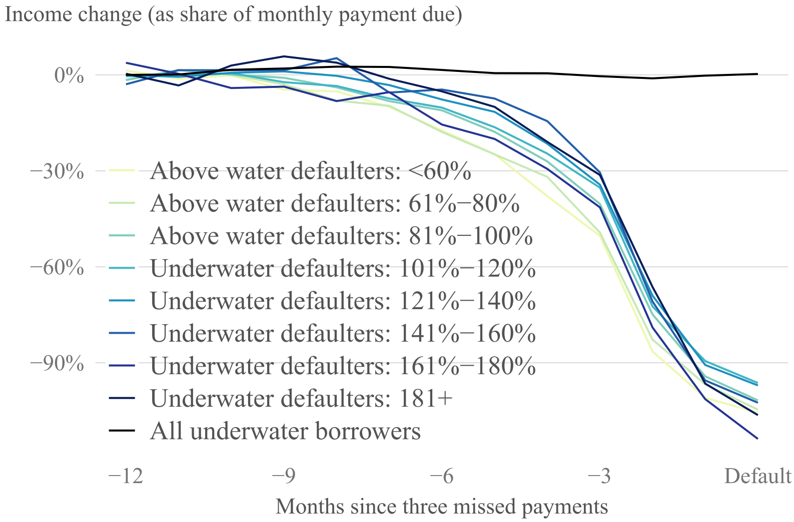

FIGURE IV

Heterogeneity in Income Drop by Loan-to-Value Ratio

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 IV.C for details.

for LTV bins in increments of 20% 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% for defaulters with LTVs between 200 and 220 and 13% for those with LTVs over 220. However, these groups account for a very small share of defaults: only 0.5% of defaulters have LTVs above 200.[22]

Figure V 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% of all underwater defaults are strategic, and that this share is rising steeply

22. Online Appendix Figure A-9 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% of defaulters had LTV > 200. This degree of negative equity requires price declines so large that this was rare, even during the Great Recession.

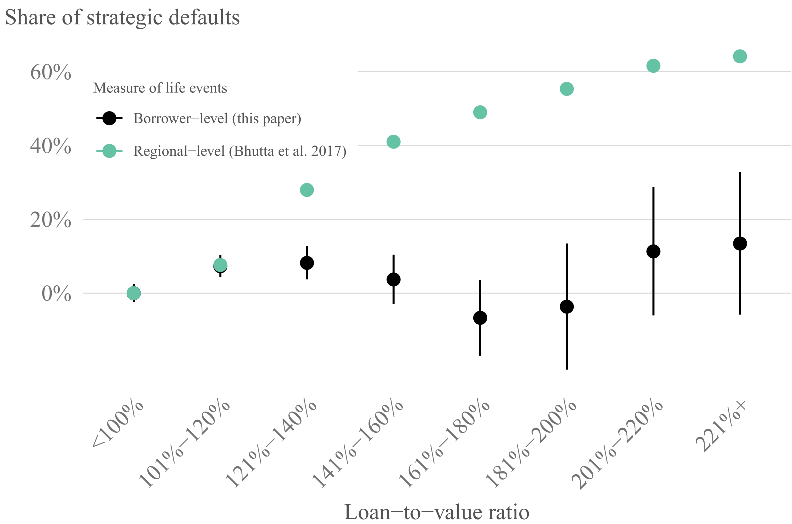

FIGURE V

Share of Strategic Defaults by Loan-to-Value Ratio

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 (7). The Bhutta, Dokko, and Shan (2017) estimates come from their figure 6, 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% confidence intervals for our estimates.

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.[23]

8. *Measurement Error in LTV Ratio*. Finally, our estimates change little after accounting for measurement error in LTVs

23. BDS says that if after controlling for regional covariates “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” (Bhutta, Dokko, and Shan 2017, 2436). 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.

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 Online Appendix 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 II must be deliberately agnostic about 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 Online Appendix Figure A-10a. 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 example, a borrower whose observed LTV is 60 has a 5% chance of being actually underwater, and a borrower whose observed LTV is 140 has a 92% chance of being actually underwater. In contrast, a borrower with observed LTV of 100 has a roughly 50% chance of being underwater.

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

$$\begin{aligned}\frac{Income_t}{\bar{Payment}_{pre}} &= \lambda_1 \mathbf{1}(G^* = 0) + \lambda_2 \mathbf{1}(G^* = 1) \\ &\quad + \lambda_3 \mathbf{1}(t = -2, -1, 0) \\ &\quad \times \mathbf{1}(G^* = 0) + \lambda_4 \mathbf{1}(t = -2, -1, 0) \times \mathbf{1}(G^* = 1) + \varepsilon. \qquad (11)\end{aligned}$$

Because we do not observe G^*, we use an adjustment procedure akin to two-sample IV, where the first-stage sample is from

CoreLogic and the second-stage sample is from JPMCI. Using $\hat{P}(G^* = 1|LTV)$ as constructed in the previous paragraph, we can feasibly estimate

$$\frac{Income_t}{\bar{Payment}_{pre}} = \lambda_1 \hat{P}(G^* = 0|LTV) + \lambda_2 \hat{P}(G^* = 1|LTV) + \lambda_3 \mathbf{1}(t = -2, -1, 0) \times \hat{P}(G^* = 0|LTV) + \lambda_4 \mathbf{1}(t = -2, -1, 0) \times \hat{P}(G^* = 1|LTV) + \varepsilon. \quad (12)$$

We calculate that the income drop before default for the truly above water is 92.5% (compared with 92.8% for those observed above water) and for the truly underwater is 88.0% (compared with 87.1% for those observed underwater).[24] Then, to convert the regression estimates to an estimate of the role of life events, we use equation (7): $\hat{\alpha}_{\text{life event}} = \frac{\hat{\lambda}_4 - \hat{\varphi}}{\hat{\lambda}_3 - \hat{\varphi}}$. Table VI, Panel A shows that $\hat{\alpha}_{\text{life event}}$ rises from 0.94 in our main specification to 0.95 when adjusting for this issue.[25] The table also shows that our estimates are similar even when we consider more extreme adjustments such as using the distribution of LTV errors from foreclosure sales; this adjustment is likely to overstate the extent of measurement error in LTV because most defaulters do not experience a completed foreclosure.

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 IV. 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 drops are 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

24. We report measurement error–corrected estimates analogous to equation (9) in Table V, column (5).

25. Thus, 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.

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.

IV.D. 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 subpopulations where we might expect this behavior to be more prominent. In fact, our method does detect more evidence of strategic behavior in two subpopulations: borrowers who miss three straight mortgage payments and subprime borrowers. 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.[26] Therefore, if there is evidence of strategic default, it should manifest in this subpopulation. 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 subgroup. Online Appendix Figure A-11a shows larger income declines for above-water defaulters than for underwater defaulters in the subsample that misses three consecutive mortgage payments. In Table VI, Panel B, we estimate that 13% of underwater defaults are strategic for this subsample.

We also study the default behavior of subprime borrowers, who account for a small share of originations (Online Appendix

26. Other papers that study consecutive missed payments include Experian and Wyman (2009); Tirupattur, Chang, and Egan (2010); Keys et al. (2013) and Bradley, Cutts, and Liu (2015). Keys et al. (2013) 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. (2013) through 2015 using the CRISM data and show the results in Online Appendix Figure A-12. We find that 15.9% of underwater defaults meet their definition of sequential missed payments, while 9.7% of above-water defaults meet this definition. The excess sequential default rate for underwater borrowers is 6.2%. If we interpret this as an alternative estimate of the prevalence of strategic default, it falls within the confidence interval of our central estimate.

Table A-1a) but a disproportionately large share of defaults (Online Appendix Table A-1b) and are sometimes hypothesized to be particularly prone to strategically default (Geanakoplos 2014). Online Appendix Figure A-11b shows larger income declines for above-water subprime defaulters than for underwater subprime defaulters and in Table VI, Panel B, we estimate that 9% of underwater defaults are strategic for this subsample. This is slightly larger than the strategic default rate in the overall population, although the difference is not statistically significant.

Finally, one more group that we might expect to disproportionately default strategically is investors (Albanesi, De Giorgi, and Nosal 2017). Although our sample is representative relative to external benchmarks in terms of the share of self-declared investors (see Online Appendix 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 Chase 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 government-sponsored enterprise, portfolio, and privately securitized loans. Moreover, because this subpopulation 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) report that fraudulent investors account for only 10% of all defaults.

IV.E. Causal Impact of Negative Equity

1. *Research Design.* We 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 cyclicality. Guren et al. (2021) further develop this instrument and uses it to estimate housing wealth elasticities.

Below, we apply the cyclicality method to estimate the causal impact of negative equity. We find that the estimates for subprime borrowers in Palmer (2015) extend to a broad sample of borrowers using Chase and CRISM data.

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

$$h_{i,t,m,j,r} = h_0(t)\exp\left\{\psi_j + \xi_{r,m} + \delta\mathbf{1}(LTV_{i,t} > 100) + X'_{i,m,j}\theta + v_{i,t,m,j,r}\right\}, \tag{13}$$

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 (e.g., February 2010), core-based statistical area (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 one if borrower i has negative equity in month t, $X'_{i,m,j}$ is a vector of controls including borrower and loan characteristics measured at origination, and $v_{i,t,m,j,r}$ is the error term.[27]

The coefficient of interest in equation (13) 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 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 IV approach.

We use the cyclicality instrument from Palmer (2015) and Guren et al. (2021). A robust empirical fact is that house prices in some cities are systematically more sensitive to regional price

27. Borrower and loan characteristics include an indicator variable for whether the loan is a balloon mortgage, 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. Below we discuss CBSA-level control variables that are also included.

cycles than are house prices in other cities, presumably because of permanent differences in elasticities of housing supply. Building new houses requires land, permits, materials, and labor. Supply elasticities can differ because of differences in land availability, the regulatory environment for permits, and the local cost of labor and materials for construction. The prior literature has taken two approaches to measuring supply elasticities: identifying observable proxies for these inputs (Saiz 2010 for land; Ganong and Shoag 2017 for permits) and constructing omnibus elasticity measures (Palmer 2015; Guren et al. 2021). The strength of the second approach is that—rather than requiring a researcher to collect information on all the different observable determinants of the housing supply elasticity (an approach that we have already seen is quite difficult in the context of negative life events)—it instead infers the combined importance of both observable and unobservable determinants of housing supply elasticity from the net effect of shocks to regional housing demand on city-level prices. 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). We describe the first-stage relationship between the instrument and LTV and then discuss the conditions under which this approach identifies the causal effect of negative equity in our setting.

The city-level cyclicality instrument is highly predictive of LTVs, indicating a strong first stage. We regress LTV on the same fixed effects and controls as in equation (13) as well as the interaction between the sensitivity measure and house price changes, that is,

$$LTV_{i,t,m,j,r} = \Psi_j + \Xi_{r,m} + \rho\Gamma_j \Delta P_{r,m} + X'_{i,m,j}\Theta + \eta_{i,t,m,j,r}, \tag{14}$$

where Γ_j is the city-level house price sensitivity measure reported in Guren et al. (2021), $\Delta P_{r,m}$ is the log annual change in regional house prices, and $\eta_{i,t,m,j,r}$ is the error term.[28] The first stage is presented visually in Online Appendix Figure A-13 using a residualized binscatter. The F-statistic is high (over 80). Intuitively,

28. Following 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 that interacts Γ_j with calendar-month fixed effects rather than regional house price changes, which we discuss below.

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 nonlinear hazard model setting of equation (13), following Palmer (2015) and Imbens and Wooldridge (2007). This involves adding the estimated residuals $\hat{\eta}_{i,t,m,j,r}$ from equation (14) as controls when estimating equation (13). The key identification assumption in this setting is that the instrument is independent of the error terms in equations (13) and (14), conditional on all the included controls. This assumption implies the two standard types of restrictions necessary for identification in an IV setting. We provide an intuitive discussion of these restrictions here and describe the restrictions formally in Online Appendix C.7.

First, the instrument $\Gamma_j \Delta P_{r,\,m}$ must be as good as randomly assigned, conditional on the controls. Without the inclusion of any controls, this restriction would require that borrowers in high-sensitivity areas are not inherently more or less likely to default than borrowers in low-sensitivity areas, regardless of negative-equity status. The inclusion of controls helps address many of the ways this assumption could fail. For example, Davidoff (2016) has critiqued the use of sensitivity instruments in regressions using a single cross section because housing supply elasticity is correlated with measures of permanent long-run demand growth. However, our setting allows us to include CBSA fixed effects, which eliminates permanent differences between CBSAs as a source of identification. Including CBSA fixed effects also addresses concerns related to systematic differences in income levels or differences in long-run income growth between high- and low-sensitivity areas.

Another concern about how the as-good-as-random assumption might fail is that workers in high-sensitivity areas could be concentrated in certain industries (Howard and Liebersohn 2020) that are differentially sensitive to recessions. For example, suppose that high-sensitivity areas also have a high share of employment in durables-adjacent industries. When a (typical) national recession hits, demand will fall for durables and employment will fall most in durables-adjacent industries. Then defaults will rise most in those places even if house prices do not change at all. More generally, any correlation in employment cyclicality and house price cyclicality may lead to a violation of the as-good-as-random assumption.

We follow Guren et al. (2021) in addressing this concern using two strategies. First, we control for two-digit local industry shares with separate coefficients for each time period. This accounts for differential factors that affect default and are correlated with industrial structure in the cross section. Second, there may be factors beyond industrial structure that cause some cities to be more sensitive to business cycles in a way that is correlated with house price sensitivity. Guren et al. (2021) develop a measure of each city's sensitivity to region-level employment cycles, and we follow them by including this measure as a control.[29]

With these controls included, the condition needed for the as-good-as-random restriction to hold is that, conditional on these controls, there is no other factor that is both correlated with regional house prices in the time series ($\Delta P_{r,m}$) and that differentially affects default risk in the high-sensitivity CBSAs (Γ_j). Although we believe this condition is plausible, we think that if it fails and such a factor actually does exist, the correlation is likely negative (i.e., when house prices fall in the time series, default risk is more likely to rise than to fall coincidentally in the high-sensitivity CBSAs). If so, then this would lead us to overestimate the causal effect of negative equity on default and therefore lead to a conservative estimate of the share of pure cash flow defaults (which is $1 - \alpha_{\text{negative equity}}$).

The second restriction implied by the identification assumption is an exclusion restriction: price changes only affect default through their effect on negative equity. This rules out other causal channels by which price changes could affect default (even if house price sensitivity were randomly assigned across CBSAs). This implies that sensitivity cannot directly affect the probability of default or the probability of a borrower experiencing a negative life event (which itself directly affects default).[30] This restriction could fail if declines in house prices in more sensitive areas

29. See equation (2) in Guren et al. (2021) for more details.

30. It is not a violation of the exclusion restriction in Online Appendix C.7 if a borrower's negative equity itself has a direct effect on that borrower's probability of a negative life event. One source of such an effect is constrained mobility (Gopalan et al. 2021). Because a borrower with negative equity cannot easily sell their house, they are less able to pursue alternative employment opportunities in different cities and so are more exposed to relative income declines. This is part of the causal effect of negative equity that our design seeks to capture. Similarly, if negative equity reduces borrower liquidity, making them less able to withstand a given income shock, this is also part of the causal effect of negative equity because it is a force that could drive higher defaults for negative equity than for positive equity borrowers even when facing the same set of negative life events.

translate into local spending declines (Mian and Sufi 2014; Aladangady 2017), which then lead to local employment declines that themselves trigger default. In principle, one potential solution to this identification threat would be to control for all observable negative life events. However, as we highlight, negative life events are inherently difficult to measure, so we do not view this as a viable solution.

We instead document results from three alternative approaches to address possible concerns about this restriction. First, the exclusion restriction should be less controversial for subsamples with little labor market income risk, and we therefore analyze one such subsample: retirees with a fixed income. If local house price declines lead to reduced spending and employment, the exclusion restriction should still hold for retirees who are less exposed to labor market risk because they are receiving a fixed income from Social Security. We therefore reestimate equation (13) for this subsample. Second, as with potential failures of the random assignment restriction, we note that failures of the exclusion restriction are likely to lead us to overestimate the causal effect of negative equity on default since, if anything, area-level house price declines are likely to increase the probability that borrowers in those areas suffer negative life events. This is therefore another reason our estimate of pure cash flow defaults may be conservative for our sample. Third, we show that an alternative research design for estimating the causal effect of negative equity on default in prior work, which does not require this restriction, leads to a similar conclusion.

2. *Causal Impact of Negative Equity on Default: Estimate and Robustness.* Table VII shows our results. Panel A shows estimates of the default hazard function in equation (13) for borrowers in the Chase sample. Column (1) shows the most parsimonious specification, without controls and without instrumenting for negative equity. The coefficient on the underwater dummy shows an estimate of $\hat{\delta} = 1.50$. Controlling for borrower and loan characteristics as well as CBSA-level industry shares and employment sensitivity in column (2) reduces this estimate slightly to 1.32.

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% ($exp(\hat{\delta})$). The statistical significance of the LTV residuals suggests that the naive underwater

TABLE VII

Impact of Negative Equity on Default

	(1)	(2)	(3)	(4)
Panel A: Chase sample				
Underwater	1.501	1.323	0.359	0.330
	(0.051)	(0.058)	(0.063)	(0.042)
LTV fitted residuals			2.078	1.942
			(0.143)	(0.108)
$\hat{\alpha}_{\text{negative equity}}$	0.777	0.734	0.302	0.281
	(0.011)	(0.015)	(0.044)	(0.03)
Region-year FEs	Y	Y	Y	N
CBSA FEs	Y	Y	Y	Y
Borrower and loan characteristics	N	Y	Y	Y
CBSA controls	N	Y	Y	Y
Origination year FEs	N	N	N	Y
Instrument	None	None	Cyclicality-HPI	Cyclicality-month
First-stage partial F-stat	—	—	81.64	16.95
Log likelihood	−340,710	−337,360	−334,187	−333,148
Observations	1,432,248	1,432,248	1,432,248	1,432,248
Panel B: CRISM sample				
Underwater	1.400	0.963	0.282	0.257
	(0.035)	(0.039)	(0.044)	(0.038)
LTV fitted residuals			1.626	1.553
			(0.077)	(0.073)
$\hat{\alpha}_{\text{negative equity}}$	0.754	0.618	0.246	0.227
	(0.009)	(0.015)	(0.033)	(0.030)
Region-year FEs	Y	Y	Y	N
CBSA FEs	Y	Y	Y	Y
Borrower and loan characteristics	N	Y	Y	Y
CBSA controls	N	Y	Y	Y
Origination year FEs	N	N	N	Y
Instrument	—	—	Cyclicality-HPI	Cyclicality-month
First-stage partial F-stat	—	—	667.76	127.7
Log likelihood	−467,501	−457,570	−454,895	−454,647
Observations	1,456,127	1,456,127	1,456,127	1,456,127

Notes: This table reports maximum-likelihood estimates of the default hazard model from equation (13). 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 $LTV > 100$. Columns (3) and (4) instrument for Underwater using the control function approach by including the LTV fitted residuals from equation (14). 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 house 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 (13)), 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, that is, $\hat{\alpha}_{\text{negative equity}}$. Denoting the coefficient on Underwater as $\hat{\delta}$ from equation (13), then $\hat{\alpha}_{\text{negative equity}} = 1 - exp(-\hat{\delta})$. See Section IV.E.

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 $\hat{\delta} = 0.33$.

Our estimates are similar when we use a different data source and are also similar to prior estimates for specific subpopulations. Panel B replicates the specifications from Panel A in the CRISM data set, which captures mortgages from a broad range of lenders. We find IV estimates of $\hat{\delta}$ between 0.26 and 0.28 in columns (3) and (4). The estimates for borrowers from both panels in Table VII (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.

Our subsample analysis suggests that if anything, the true causal effect of negative equity may be even smaller than these estimates suggest. Online Appendix Table A-6 shows $\hat{\delta}$ estimates of 0.12 to 0.15 for retirees, who may be especially likely to meet the identification assumptions because their incomes are less likely to be correlated with house price declines.[31] However, this subsample has a lower default rate and a lower negative equity share than the full population, so the smaller causal estimate could also be due to different subsample characteristics.

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%. Combining the causal effect of negative equity with the formula from equation (4), $\alpha_{\text{negative equity}} = 1 - E\left[\frac{Y(T^*,1)}{Y(T^*,0)}|G = 1\right]^{-1} = 1 - exp(-\delta)$. At our preferred estimate of $\hat{\delta} = 0.36$, $\hat{\alpha}_{\text{negative equity}} = 30\%$. Estimates of $\hat{\alpha}_{\text{negative equity}}$ from each specification are reported in the bottom rows of Table VII, along with standard errors calculated using the delta method. The IV estimates are similar

31. Online Appendix Table A-7 shows no significant correlation between the instrument and observable measures of individual income (both current and future bank account income). This is true both in the retiree sample (Panel B) and in the full sample (Panel A). Although these patterns support the identification assumptions for both samples, the retiree sample is still useful because there could be a correlation between the instrument and unobservable measures of income, and such a correlation is a priori less likely for retirees.

across two different ways of constructing the instrument and two different data sets, ranging from 21% to 30%.

We would reach a similar conclusion if, rather than directly estimating the causal effect of negative equity on default and relying on the identification assumptions necessary in our setting, we combined our framework with an alternative research design in prior work. Gupta and Hansman (2022) estimates the causal effect 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 individual-level exogenous variation in mortgage balances between otherwise similar borrowers. It therefore does not rely on a regional instrument for negative equity and does not need the identifying assumptions discussed above (though of course it relies on alternative assumptions). The paper finds that negative equity causally increases default probabilities by about 30 percentage points and that the underwater default probability is about 60%. Plugging these numbers into equation (4) delivers an estimate of $\hat{\alpha}_{\text{negative equity}} = 50\%$ for their sample of option ARM borrowers.

IV.F. Decomposition: Strategic versus Cash Flow versus Double-Trigger Defaults

Using $\hat{\alpha}_{\text{life event}}$ from Section IV.B and $\hat{\alpha}_{\text{negative equity}}$ from Section IV.E 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% of underwater defaults. Interpreted through the lens of the potential-outcomes model in Table II, this suggests that 30% 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% of underwater defaults are cash flow defaults. Combined with our prior estimate from $\hat{\alpha}_{\text{life event}}$ that 6% of underwater defaults are strategic, this means 24% 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.[32] 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, which commonly assumes that any underwater default that is not strategic must be double trigger.[33] Our preferred estimate is that the commonly neglected cash flow defaults account for 70% of all underwater defaults. Two alternative estimates also point to a meaningful fraction of pure cash flow defaults. First, instead of using our measure of $\hat{\alpha}_{\text{negative equity}}$ estimated for our full sample of borrowers, we could apply the estimate from Gupta and Hansman (2022). Even among their sample of option ARM borrowers, who they note are likely to be more sophisticated (and therefore may be more attuned to the financial benefits of exercising their default option), the 50% estimate of $\hat{\alpha}_{\text{negative equity}}$ implies that the remaining 50% of defaults among these borrowers are driven purely by cash flow. Second, even if one were to take the extreme assumption that the entire cross-sectional relationship between negative equity and default is causal, one would still find that a significant fraction of defaults are driven entirely by cash flow. This specification, from Table VII column (1), implies that the pure cash flow channel accounts for about one-quarter of all underwater defaults.

Furthermore, two robustness analyses suggest that the true share of pure cash flow defaults could be even larger than our preferred estimate of 70%. In the retiree subsample, the share of defaults caused entirely by cash flow is 86% to 88%. In addition, if we consider an alternative thought experiment of reassigning underwater borrowers to an LTV of 90 (rather than to the average behavior of all above-water borrowers), the share of defaults

32. See also Cunningham, Gerardi, and Shen (2021) for more evidence of double-trigger behavior.

33. See 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.

caused entirely by cash flow is 81% to 86% (Online Appendix Table A-8).

It is important to emphasize that our finding of little strategic default stands alongside (rather than in tension with) the prior literature's finding that underwater borrowers are more likely to default than above-water borrowers. The strong cross-sectional relationship between negative equity and default that has been documented in the prior literature also holds in our data set (Table III). How can both findings be true? This can arise either because underwater borrowers are more likely to experience negative life events (cash flow default) or because underwater borrowers are more likely to default conditional on a life event (double-trigger default). Our decomposition between the three theories of default suggests that both factors play an important role.

V. Reanalysis of Survey Data

Our empirical findings may be surprising because some prior work estimates that between 30% and 70% of Great Recession defaults were strategic (Guiso, Sapienza, and Zingales 2013; Bhutta, Dokko, and Shan 2017; Gerardi et al. 2018). It is natural to wonder whether our 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 helps address a concern about whether underwater borrowers hide some of their income from the bank that services their mortgage.[34] Second, it captures borrowers' perceived LTV, which is the decision-relevant measure of LTV

34. This 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

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) survey a representative sample of U.S. households, about one-third of whom report that they know someone who has defaulted on their mortgage. On page 1474, 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 analyze the PSID because it includes information on whether borrowers are above- or underwater, which allows us to implement our comparison group approach.

To measure mortgage affordability empirically in the PSID, GHOW examine the distribution of disposable income (income y minus nonhousing 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" $= y - (c + m) \geqslant 0$) and they nevertheless default, they must be defaulting strategically.

We reanalyze the prevalence of underwater strategic default in the 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 article, but directly applying the method from Section II yields statistically imprecise results.[35]

---

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.

35. Our reverse-regression 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 side of a regression creates larger standard errors, but not bias. Our method, which places a noisy measure of treatment on the left side, sacrifices precision to avoid bias. This is not a problem for a large administrative data set, but it is a problem for the PSID, where there are only 244 mortgage defaulters.

Instead, we use the can pay definition of strategic default to enable comparability to the prior literature.

Figure VI, Panel A plots the cumulative distribution function of available resources for above-water defaulters, underwater defaulters, and all underwater borrowers in the PSID. Figure VI, Panel B constructs the analogous affordability statistic in the JPMCI data. Specifically we measure y in the quarter of default and $c + m$ as bank account expenditures three quarters prior to default. Consistent with a role for negative life events, available resources are much lower for defaulters than for nondefaulters. In contrast, when we look within defaulters to compare above- and underwater borrowers, the distribution of available resources appears to be the same. Thus, across both the PSID and JPMCI, the GHOW affordability statistic shows a similar degree of distress for above- and underwater defaulters—a finding that suggests little strategic default.

The key reason we find little strategic default when prior work found substantial strategic default is that our estimation method adjusts for measurement error, rather than differences in data sources or differences in the definition of strategic default. For example, Online Appendix Figure A-14a shows that 39% (33.3%) of PSID (JPMCI) underwater defaulters meet the can pay definition, and one might then conclude that these 33%–39% of underwater defaults are strategic. Yet Online Appendix Figure A-14a shows that 37% (33.6%) of PSID (JPMCI) above-water defaulters also meet the can pay definition in a sample that has no motive to default strategically. The difference in the share of defaults labeled can pay for above-water versus underwater is not statistically significant in either data set. We conclude that income and consumption obligations are difficult to measure at the household level.[36] 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

36. For example, one challenge for survey-based evidence is that 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). In Section III we discuss some of the sources of measurement error in administrative bank data.

(A) Panel Study of Income Dynamics

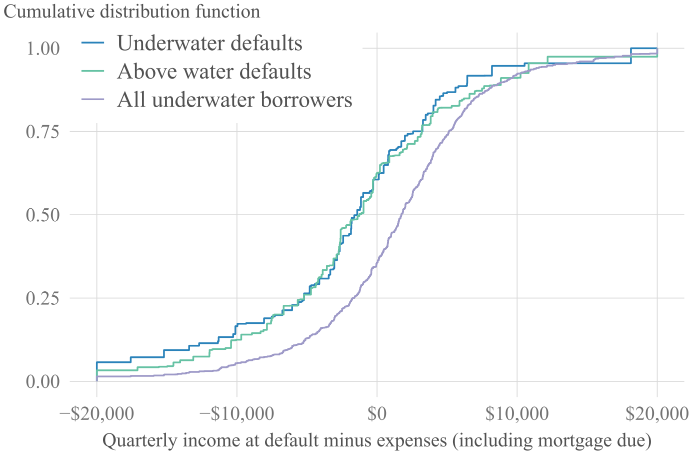

(B) Bank Account Data

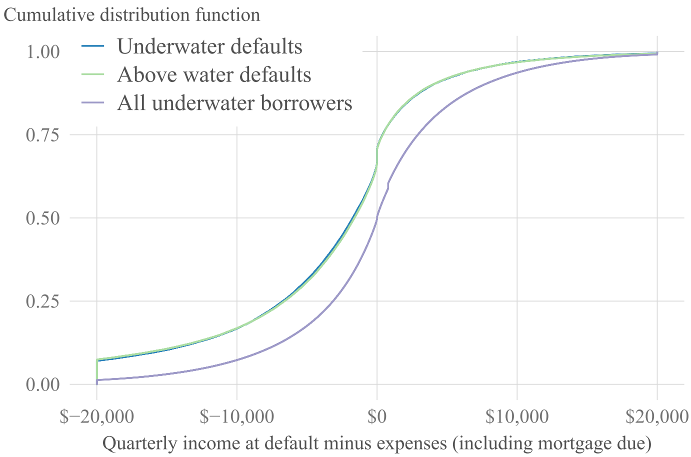

FIGURE VI

The Distribution of Available Resources is Similar for Above-Water and Underwater Defaulters

This figure shows the cumulative distribution function of available resources by home equity and default status in the Panel Study of Income Dynamics (PSID) and the bank account data. We follow Gerardi et al. (2018) in defining available resources as quarterly income minus nonhousing consumption and mortgage payment due and defining default as 60+ days past due. We use their exact definition in the PSID, and we construct the analogue in the bank account data using current income minus lagged expenses (both housing and nonhousing). We winsorize this variable at ± $20,000. The figure shows the distribution separately for above-water defaulters, underwater defaulters, and all underwater borrowers. See Section V for details.

defaulters rather than a result of one specific definition of mortgage affordability. GHOW also analyzes an alternative definition of mortgage affordability in the PSID, which examines how many people would "need to reduce consumption below subsistence levels to remain current on their mortgage." Using this measure of affordability, Online Appendix Figure A-14a shows that the share of defaults labeled as strategic is quantitatively similar for above- and underwater borrowers (p-value of .23). Furthermore, Online Appendix Figure A-15a shows that the entire distribution of $y - (c_{subsistence} + m)$ is similar. Because above- and underwater defaulters have the same distribution of available resources across different measures of mortgage affordability, we conclude that there is little evidence of strategic default in the PSID.[37]

Finally, looking beyond the PSID, in subsequent work Low (2021) studies the reasons for mortgage default using the American Survey of Mortgage Borrowers. This new survey was specifically designed to capture as many life events as possible. The paper finds that most defaulters have experienced an adverse life event and estimates that at most 4% of defaults are strategic. This similar finding from an alternative data set provides further evidence that our results are not driven by our particular analysis sample.

VI. 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 into 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

37. These figures define underwater as LTV greater than 100, to be consistent with the rest of our article. Online Appendix Figures A-14b and A-15b present the same analysis using an LTV cutoff of 90, which is the cutoff used in GHOW.

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 (Foster and Van Order 1984; Epperson et al. 1985; Deng, Quigley, and Van Order 2000). CC adds idiosyncratic income risk to the option value model to quantify the relative contributions of negative equity and cash flow motives to the default decision. Relative to 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 II we assumed that each treatment was binary (i.e., 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 V 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^{*}.[38]

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.[39] The figure shows that among households with LTV between 100 and 120, the income of defaulters is substantially lower than the income of nondefaulters. On page 1496, CC explains that this pattern emerges because short-term cash flow considerations drive the default decisions of moderately underwater borrowers: "As

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

39. This figure evaluates the joint distribution of income, home equity, and default for mortgagors with ARMs. 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. Online Appendix Figure A-16 shows that we can replicate the summary statistics from CC's Table II, panel D.

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 nondefaulters. 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 in the model. This requires two extensions. First, instead of comparing the level of income of defaulters to that of nondefaulters, our approach relies on calculating the within-borrower change in income prior to default. We show in Figure VII, Panel A how this statistic varies with LTV both in the baseline CC model and in our data.[40] 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 VII, Panel A 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 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.

40. The time interval in the CC 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 capture the difference in income around default relative to its stable baseline level prior to default.

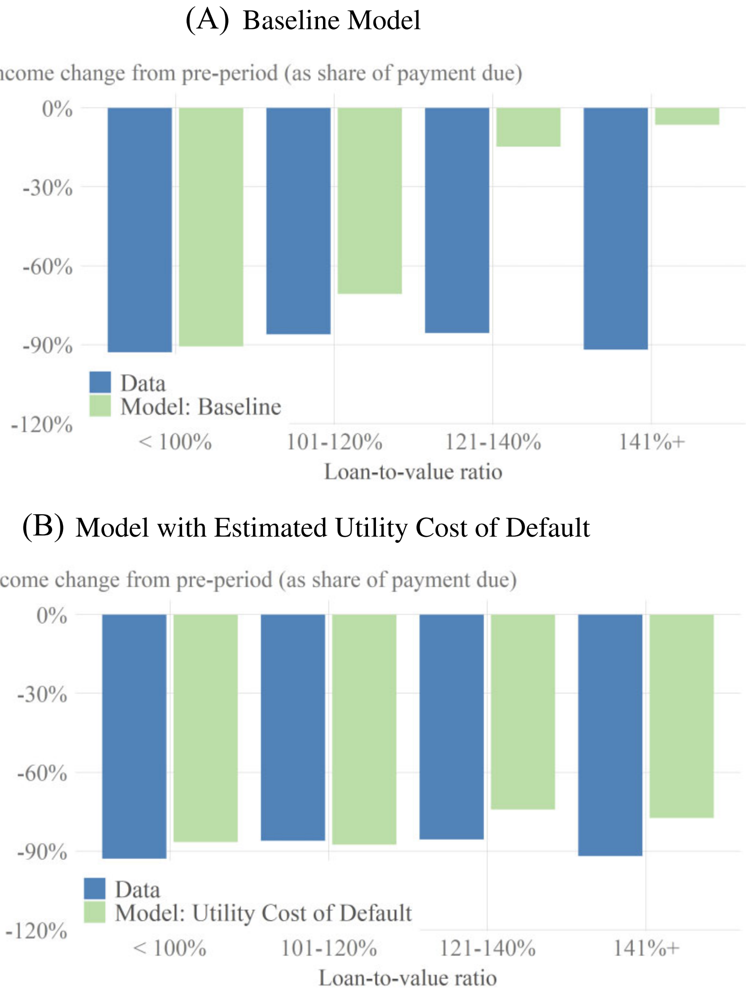

FIGURE VII

The Income Drop at Default Compared to Prior Theoretical Literature

This figure compares the 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% 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 VI for details.

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 VII, Panel A 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 IV, 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 VII, Panel A shows that the model predicts a drop in income (as a share of mortgage payment due) of 15% 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 defaulting borrowers have mean liquid assets equal to slightly more than one year of income. This contrasts with the finding in Table IV that underwater checking account balances prior to default are equal to less than two weeks of income, just like the above-water borrowers. Online Appendix Table A-9 further shows that at the time of default, in the data the entire distribution of checking account balances of underwater borrowers is similar to that of above-water 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 call "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 nonfinancial costs due to borrowers' attachment to their current home (i.e., their idiosyncratic private valuation of the home may be greater than the market's valuation), 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, their 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% decrease in the constant-equivalent consumption stream to as high as a 70% decrease (Hembre 2018; Laufer 2018; Schelkle 2018; Kaplan, Mitman, and Violante 2020).

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% decrease in the constant-equivalent consumption stream, or $100,000 in present-value terms (see Online Appendix E 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% 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) and that borrowers perceive a high cost of forced displacement above and beyond the financial and moral costs of default (Collier, Ellis, and Keys 2021). Understanding the exact sources of this cost is outside the scope of this article, but we think it is an important topic for future research.

Under this alternative parameterization, the model is able to closely match our new empirical moments. This is shown visually in Figure VII, Panel B. 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 microfoundation for the behavior we observe empirically in 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.[41] 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 Online Appendix Tables A-9 and A-10, we provide moments of 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 life events into macroeconomic models is an exciting topic for future work.

41. For example, see Corbae and Quintin (2015); Mitman (2016); Garriga and Hedlund (2020); Kaplan, Mitman, and Violante (2020); Campbell, Clara, and Cocco (2021); Greenwald, Landvoigt, and Van Nieuwerburgh (2021); Guren, Krishnamurthy, and Mcquade (2021); Chodorow-Reich, Guren, and McQuade (2022); and Diamond and Landvoigt (2022).

VII. Conclusion

This article 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 long-standing interest, it has remained difficult to distinguish between these triggers in part because it is difficult to precisely measure life events. We address this measurement error problem using a comparison group of borrowers whose default must have been caused by a negative life event: borrowers with positive home equity. For these borrowers, negative equity cannot be the cause of their default. We implement this method in a new high-frequency data set linking income and mortgage default.

Our central finding is that only 6% of underwater defaults are caused exclusively by negative equity, much less than previously thought. Furthermore, we find that 70% of underwater defaults are driven solely by negative life events. Moreover, because even at the peak of the housing crisis at least 40% 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 macro-finance models where borrower default decisions play a central role.

One interesting avenue for future work would be to develop models that can match all of the available evidence on the joint variation of income, home equity, and mortgage default. We show that a high utility cost of mortgage default is sufficient to match our core empirical finding that even significantly underwater borrowers do not default on their mortgages unless they have also suffered a negative life event. However, richer models may be needed to match additional aspects of the available evidence. Three examples demonstrate why this might be useful. First, a higher utility cost of default lowers overall default rates, so other model parameters may need to adjust to match aggregate default rates in the data. Second, matching the fact that default rates rise gradually with LTV may require adding in a correlation between LTV and negative life events. Third, it would be useful to microfound high default costs, ideally in conjunction with additional empirical work disentangling the sources of these costs.

An additional contribution of this article is that it may be helpful when 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 policy makers 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 immediate liquidity while leaving the principal balance unchanged (Campbell, Clara, and Cocco 2021).

This logic has support in the prior literature, which, using within-contract variation among borrowers who received a mortgage modification, has shown that reducing payments is more effective than forgiving principal at preventing default (Ganong and Noel 2020; Sharlemann and Shore 2022). 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 study 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 offered immediate liquidity via forbearance (Cherry et al. 2021). If house prices fall again and unemployment rises, 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.

University of Chicago, United States
University of Chicago, United States

Supplementary Material

An Online Appendix for this article can be found at *The Quarterly Journal of Economics* online.

Data Availability

Data and code replicating the tables and figures in this article can be found in Ganong and Noel (2022) in the Harvard Dataverse, https://doi.org/10.7910/DVN/SDNMSR.

References

Aladangady, Aditya, "Housing Wealth and Consumption: Evidence from Geographically Linked Microdata," *American Economic Review*, 107 (2017), 3415–3446. http://dx.doi.org/10.1257/aer.20150491

Albanesi, Stefania, Giacomo De Giorgi, and Jaromir Nosal, "Credit Growth and the Financial Crisis: A New Narrative," NBER Working Paper no. 23740, 2017. https://doi.org/10.3386/w23740

Bayer, Patrick, Fernando Ferreira, and Stephen L. Ross, "The Vulnerability of Minority Homeowners in the Housing Boom and Bust," *American Economic Journal: Economic Policy*, 8 (2016), 1–27. http://dx.doi.org/10.1257/pol.20140074

Bernanke, Ben S., "Mortgage Delinquencies and Foreclosures," speech at the Columbia Business School's 32nd annual dinner, New York, 2008.

Bernstein, Asaf, "Negative Home Equity and Household Labor Supply," *Journal of Finance*, 76 (2021), 2963–2995. http://dx.doi.org/10.1111/jofi.13070

Bhutta, Neil, Jane Dokko, and Hui Shan, "Consumer Ruthlessness and Mortgage Default during the 2007 to 2009 Housing Bust," *Journal of Finance*, 72 (2017), 2433–2466. http://dx.doi.org/10.1111/jofi.12523

Boar, Corina, Denis Gorea, and Virgiliu Midrigan, "Liquidity Constraints in the U.S. Housing Market," *Review of Economic Studies*, 89 (2022), 1120–1154. https://doi.org/10.1093/restud/rdab063

Bradley, Michael G., Amy Crews Cutts, and Wei Liu, "Strategic Mortgage Default: The Effect of Neighborhood Factors," *Real Estate Economics*, 43 (2015), 271–299. http://dx.doi.org/10.1111/1540-6229.12081

Campbell, John Y., Nuno Clara, and João F. Cocco, "Structuring Mortgages for Macroeconomic Stability," *Journal of Finance*, 76 (2021), 2525–2576. http://dx.doi.org/10.1111/jofi.13056

Campbell, John Y., and João F. Cocco, "A Model of Mortgage Default," *Journal of Finance*, 70 (2015), 1495–1554. http://dx.doi.org/10.1111/jofi.12252

Campbell, Tim S., and J. Kimball Dietrich, "The Determinants of Default on Insured Conventional Residential Mortgage Loans," *Journal of Finance*, 38 (1983), 1569–1581. http://dx.doi.org/10.1111/j.1540-6261.1983.tb03841.x

Chan, Sewin, Andrew Haughwout, Andrew Hayashi, and Wilbert Van Der Klaauw, "Determinants of Mortgage Default and Consumer Credit Use: The Effects of Foreclosure Laws and Foreclosure Delays," *Journal of Money, Credit and Banking*, 48 (2016), 393–413. http://dx.doi.org/10.1111/jmcb.12304

Cherry, Susan, Erica Xuewei Jiang, Gregor Matvos, Tomasz Piskorski, and Amit Seru, "Government and Private Household Debt Relief during COVID-19," *Brookings Papers on Economic Activity*, Fall (2021), 141–221. http://dx.doi.org/10.1353/eca.2022.0002

Chodorow-Reich, Gabriel, Adam M. Guren, and Timothy J. McQuade, "The 2000s Housing Cycle with 2020 Hindsight: A Neo-Kindlebergerian View," NBER Working Paper no. 29140, 2022. https://doi.org/10.3386/w29140

Collier, Benjamin L., Cameron Ellis, and Benjamin J. Keys, "The Cost of Consumer Collateral: Evidence from Bunching," NBER Working Paper no. 29527, 2021. https://doi.org/10.3386/w29527

Corbae, Dean, and Erwan Quintin, "Leverage and the Foreclosure Crisis," *Journal of Political Economy*, 123 (2015), 1–65. http://dx.doi.org/10.1086/677349

Cunningham, Chris, Kristopher Gerardi, and Lily Shen, "The Double Trigger for Mortgage Default: Evidence from the Fracking Boom," *Management Science*, 67 (2021), 3321–3984. http://dx.doi.org/10.1287/mnsc.2020.3647

Davidoff, Thomas, "Supply Constraints Are Not Valid Instrumental Variables for Home Prices Because They Are Correlated with Many Demand Factors," *Critical Finance Review*, 5 (2016), 177–206. http://dx.doi.org/10.1561/104.00000037

Dávila, Eduardo, "Using Elasticities to Derive Optimal Bankruptcy Exemptions," *Review of Economic Studies*, 87 (2020), 870–913. http://dx.doi.org/10.1093/restud/rdz043

DeFusco, Anthony A., and John Mondragon, "No Job, No Money, No Refi: Frictions to Refinancing in a Recession," *Journal of Finance*, 75 (2020), 2327–2376. http://dx.doi.org/10.1111/jofi.12952

Deng, Yongheng, John M. Quigley, and Robert Van Order, "Mortgage Terminations, Heterogeneity and the Exercise of Mortgage Options," *Econometrica*, 68 (2000), 275–307. http://dx.doi.org/10.1111/1468-0262.00110

Diamond, Rebecca, Adam Guren, and Rose Tan, "The Effect of Foreclosures on Homeowners, Tenants, and Landlords," NBER Working Paper no. 27358, 2020. https://doi.org/10.3386/w27358

Diamond, William, and Tim Landvoigt, "Credit Cycles with Market-Based Household Leverage," *Journal of Financial Economics*, 146 (2022), 726–753. http://dx.doi.org/10.1016/j.jfineco.2021.11.001

Elul, Ronel, Aaron Payne, and Sebatian Tilson, "Owner-Occupancy Fraud and Mortgage Performance," Federal Reserve Bank of Philadelphia Working Paper 19–53, 2021. https://doi.org/10.21799/frbp.wp.2019.53

Epperson, James F., James B. Kau, Donald C. Keenan, and Walter J. Muller, "Pricing Default Risk in Mortgages," *Real Estate Economics*, 13 (1985), 261–272. http://dx.doi.org/10.1111/1540-6229.00354

Experian and Oliver Wyman, "Experian–Oliver Wyman Market Intelligence Report: Understanding Strategic Default in Mortgages. Part I," Experian and Oliver Wyman Technical Report, 2009.

Farrell, Diana, Kanav Bhagat, Peter Ganong, and Pascal Noel, "Mortgage Modifications after the Great Recession: New Evidence and Implications for Policy," JPMorgan Chase Institute Technical Report, 2017. http://dx.doi.org/10.2139/ssrn.3083560

Farrell, Diana, Kanav Bhagat, and Chen Zhao, "Falling Behind: Bank Data on the Role of Income and Savings in Mortgage Default," JPMorgan Chase Institute Technical Report, 2018. http://dx.doi.org/10.2139/ssrn.3273062

Foote, Christopher L., Kristopher Gerardi, and Paul S. Willen, "Negative Equity and Foreclosure: Theory and Evidence," *Journal of Urban Economics*, 64 (2008), 234–245. http://dx.doi.org/10.1016/j.jue.2008.07.006

Foote, Christopher L., and Paul S. Willen, "Mortgage-Default Research and the Recent Foreclosure Crisis," *Annual Review of Financial Economics*, 10 (2018), 59–100. http://dx.doi.org/10.1146/annurev-financial-110217-022541

Foster, Chester, and Robert Van Order, "An Option-Based Model of Mortgage Default," *Housing Finance Review*, 3 (1984), 351–372.

Ganong, Peter, and Pascal Noel, "Liquidity vs. Wealth in Household Debt Obligations: Evidence from Housing Policy in the Great Recession," *American Economic Review*, 110 (2020), 3100–3138. http://dx.doi.org/10.1257/aer.20181243

———, "Replication Data for: 'Why Do Borrowers Default on Mortgages?'," (2022), Harvard Dataverse, https://doi.org/10.7910/DVN/SDNMSR.

Ganong, Peter, and Daniel Shoag, "Why Has Regional Income Convergence in the U.S. Declined?," *Journal of Urban Economics*, 102 (2017), 76–90. http://dx.doi.org/10.1016/j.jue.2017.07.002

Garriga, Carlos, and Aaron Hedlund, "Mortgage Debt, Consumption, and Illiquid Housing Markets in the Great Recession," *American Economic Review*, 110 (2020), 1603–1634. http://dx.doi.org/10.1257/aer.20170772

Geanakoplos, John, "Leverage, Default, and Forgiveness: Lessons from the American and European Crises," *Journal of Macroeconomics*, 39 (2014), 313–333. http://dx.doi.org/10.1016/j.jmacro.2014.01.001

Gerardi, Kristopher, Kyle F. Herkenhoff, Lee E. Ohanian, and Paul S. Willen, "Can't Pay or Won't Pay? Unemployment, Negative Equity, and Strategic Default," NBER Working Paper no. 21630, 2015. https://doi.org/10.3386/w21630

———, "Can't Pay or Won't Pay? Unemployment, Negative Equity, and Strategic Default," *Review of Financial Studies*, 31 (2018), 1098–1131. http://dx.doi.org/10.1093/rfs/hhx115

Ghent, Andra C., and Marianna Kudlyak, "Recourse and Residential Mortgage Default: Evidence from US States," *Review of Financial Studies*, 24 (2011), 3139–3186. http://dx.doi.org/10.1093/rfs/hhr055

Gilbukh, Sonia, and Paul S. Goldsmith-Pinkham, "Heterogeneous Real Estate Agents and the Housing Cycle," Mimeo, 2021. http://dx.doi.org/10.2139/ssrn.3436797

Goldberg, Lawrence, and Charles A. Capone, "Multifamily Mortgage Credit Risk: Lessons from Recent History," *Cityscape*, 4 (1998), 93–113. http://dx.doi.org/10.2139/ssrn.149409

Goldberger, Arthur S., "Reverse Regression and Salary Discrimination," *Journal of Human Resources*, 19 (1984), 293–318. http://dx.doi.org/10.2307/145875

Goodman, Laurie S., Roger Ashworth, Brian Landy, and Ke Yin, "Negative Equity Trumps Unemployment in Predicting Defaults," *Journal of Fixed Income*, 19 (2010), 67–72. http://dx.doi.org/10.3905/JFI.2010.19.4.067

Gopalan, Radhakrishnan, Barton H. Hamilton, Ankit Kalda, and David Sovich, "Home Equity and Labor Income: The Role of Constrained Mobility," *Review of Financial Studies*, 34 (2021), 4619–4662. http://dx.doi.org/10.1093/rfs/hhaa136

Greenwald, Daniel L., Tim Landvoigt, and Stijn Van Nieuwerburgh, "Financial Fragility with SAM?," *Journal of Finance*, 76 (2021), 651–706. http://dx.doi.org/10.1111/jofi.12992

Gross, Tal, Raymond Kluender, Feng Liu, Matthew J. Notowidigdo, and Jialan Wang, "The Economic Consequences of Bankruptcy Reform," *American Economic Review*, 111 (2021), 2309–2341. http://dx.doi.org/10.1257/aer.20191311

Guiso, Luigi, Paola Sapienza, and Luigi Zingales, "The Determinants of Attitudes toward Strategic Default on Mortgages," *Journal of Finance*, 68 (2013), 1473–1515. http://dx.doi.org/10.1111/jofi.12044

Gupta, Arpit, and Christopher Hansman, "Selection, Leverage, and Default in the Mortgage Market," *Review of Financial Studies*, 35 (2022), 720–770. https://doi.org/10.1093/rfs/hhab052

Gupta, Arpit, Edward R. Morrison, Catherine Fedorenko, and Scott Ramsey, "Leverage, Default, and Mortality: Evidence from Cancer Diagnoses," SSRN Scholarly Paper ID 2583975, 2017. https://dx.doi.org/10.2139/ssrn.2583975

Guren, Adam M., "House Price Momentum and Strategic Complementarity," *Journal of Political Economy*, 126 (2018), 1172–1218. http://dx.doi.org/10.1086/697207

Guren, Adam M., Arvind Krishnamurthy, and Timothy J. Mcquade, "Mortgage Design in an Equilibrium Model of the Housing Market," *The Journal of Finance*, 76 (2021), 113–168. http://dx.doi.org/10.1111/jofi.12963

Guren, Adam M., Alisdair McKay, Emi Nakamura, and Jón Steinsson, "Housing Wealth Effects: The Long View," *Review of Economic Studies*, 88 (2021), 669–707. http://dx.doi.org/10.1093/restud/rdaa018

Gyourko, Joseph, and Joseph Tracy, "Reconciling theory and Empirics on the Role of Unemployment in Mortgage Default," *Journal of Urban Economics*, 80 (2014), 87–96. http://dx.doi.org/10.1016/j.jue.2013.10.005

Hembre, Erik, "HAMP, Home Attachment, and Mortgage Default," University of Illinois at Chicago Working Paper, 2018.

Howard, Greg, and Jack Liebersohn, "Regional Divergence and House Prices," Charles A. Dice Working Paper 2020-04, 2020. http://dx.doi.org/10.2139/ssrn.3570654
Hsu, Joanne W., David A. Matsa, and Brian T. Melzer, "Unemployment Insurance as a Housing Market Stabilizer," *American Economic Review*, 108 (2018), 49–81. http://dx.doi.org/10.1257/aer.20140989
Imbens, Guido, and Jeffrey M. Wooldridge, "Control Function and Related Methods," *What's New in Econometrics* (2007).
Indarte, Sasha, "Moral Hazard versus Liquidity in Household Bankruptcy," SSRN working paper, 2020. http://dx.doi.org/10.2139/ssrn.3378669
Kaplan, Greg, Kurt Mitman, and Giovanni L. Violante, "The Housing Boom and Bust: Model Meets Evidence," *Journal of Political Economy*, 128 (2020), 3285–3345. http://dx.doi.org/10.1086/708816
Keys, Benjamin J., Tomasz Piskorski, Amit Seru, and Vikrant Vig, "Mortgage Financing in the Housing Boom and Bust," in *Housing and the Financial Crisis*, Edward Glaeser and Todd Sinai, eds. (Chicago: University of Chicago Press, 2013). Ch. 4, 143–204. http://dx.doi.org/10.7208/chicago/9780226030616.003.0005
Laufer, Steven, "Equity Extraction and Mortgage Default," *Review of Economic Dynamics*, 28 (2018), 1–33. http://dx.doi.org/10.1016/j.red.2017.07.006
Low, David, "Mortgage Default with Positive Equity," NYU Job Market Paper, 2018.
———, "What Triggers Mortgage Default? New Evidence from Linked Administrative and Survey Data," Consumer Financial Protection Bureau, Office of Research Working Paper no. 2, 2021.
Mayer, Christopher, Edward Morrison, Tomasz Piskorski, and Arpit Gupta, "Mortgage Modification and Strategic Behavior: Evidence from a Legal Settlement with Countrywide," *American Economic Review*, 104 (2014), 2830–2857. http://dx.doi.org/10.1257/aer.104.9.2830
Mian, Atif, and Amir Sufi, "What Explains the 2007–2009 Drop in Employment?," *Econometrica*, 82 (2014), 2197–2223. http://dx.doi.org/10.3982/ECTA10451
Mitman, Kurt, "Macroeconomic Effects of Bankruptcy and Foreclosure Policies," *American Economic Review*, 106 (2016), 2219–2255. http://dx.doi.org/10.1257/aer.20120512
O'Malley, Terry, "The Impact of Repossession Risk on Mortgage Default," *Journal of Finance*, 76 (2021), 623–650. http://dx.doi.org/10.1111/jofi.12990
Palmer, Christopher, "Why Did So Many Subprime Borrowers Default during the Crisis: Loose Credit or Plummeting Prices?," University of California, Berkeley, Working Paper, 2015. http://dx.doi.org/10.2139/ssrn.2665762
Passero, William, Thesia I. Garner, and Clinton McCully, "Understanding the Relationship: CE Survey and PCE," in *Improving the Measurement of Consumer Expenditures,* Christopher Carroll, Thomas Crossley, and John Sabelhaus, eds. (Chicago: University of Chicago Press, 2014), 181–203.
Pearl, Judea, "Probabilities of Causation: Three Counterfactual Interpretations and Their Identification," *Synthese*, 121 (1999), 93–149. http://dx.doi.org/10.1145/3501714.3501735
Piskorski, Tomasz, Amit Seru, and James Witkin, "Asset Quality Misrepresentation by Financial Intermediaries: Evidence from the RMBS Market," *Journal of Finance*, 70 (2015), 2635–2678. http://dx.doi.org/10.1111/jofi.12271
Riddiough, Timothy J., "Equilibrium Mortgage Default Pricing with Non-Optimal Borrower Behavior," PhD diss., University of Wisconsin (1991).
Riddiough, Timothy J., and Steve B. Wyatt, "Strategic Default, Workout, and Commercial Mortgage Valuation," *Journal of Real Estate Finance and Economics*, 9 (1994), 5–22. https://doi.org/10.1007/BF01153586
Rosenbaum, Paul R., "Effects Attributable to Treatment: Inference in Experiments and Observational Studies with a Discrete Pivot," *Biometrika*, 88 (2001), 219–231. http://dx.doi.org/10.1093/biomet/88.1.219
Roubini, Nouriel, "The Forthcoming 'Jingle Mail' Tsunami: 10 to 15 Million Households Likely to Walk Away from their Homes/Mortgages Leading to a Systemic Banking Crisis," *Roubini Global Economics Monitor* (2008).

Rubin, Donald B., "Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies," *Journal of Educational Psychology*, 66 (1974), 688–701. http://dx.doi.org/10.1037/h0037350

Saiz, Albert, "The Geographic Determinants of Housing Supply," *Quarterly Journal of Economics*, 125 (2010), 1253–1296. http://dx.doi.org/10.1162/qjec.2010.125.3.1253

Scharlemann, Therese C., and Stephen H. Shore, "The Effect of Changing Mortgage Payments on Default and Prepayment: Evidence from HAMP Resets," *Real Estate Economics*, 50 (2022), 1231–1256. http://dx.doi.org/10.1111/1540-6229.12377

Schelkle, Thomas, "Mortgage Default during the U.S. Mortgage Crisis," *Journal of Money, Credit and Banking*, 50 (2018), 1101–1137. http://dx.doi.org/10.1111/jmcb.12546

Tirupattur, V., O. Chang, and J. Egan, "Understanding Strategic Defaults," Morgan Stanley Research Technical Report, 2010.

Vandell, Kerry D., "How Ruthless Is Mortgage Default? A Review and Synthesis of the Evidence," *Journal of Housing Research*, 6 (1995), 245–264.

Wooldridge, Jeffrey M., *Econometric Analysis of Cross Section and Panel Data*, (Cambridge, MA: MIT Press, 2010).

Yamamoto, Teppei, "Understanding the Past: Statistical Analysis of Causal Attribution," *American Journal of Political Science*, 56 (2012), 237–256. http://dx.doi.org/10.1111/j.1540-5907.2011.00539.x