WORKING PAPER · NO. 2020-100

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

Peter Ganong and Pascal J. Noel

JULY 2020

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

Peter Ganong and Pascal J. Noel

July 2020

JEL No. E20,G21,R21

ABSTRACT

There are two prevailing theories of borrower default: strategic default—when debt is too high relative to the value of the house—and adverse life events—such that the monthly payment is too high relative to available resources. It has been challenging to test between these theories in part because adverse events are measured with error, possibly leading to attenuation bias. We develop a new method for addressing this measurement error using a comparison group of borrowers with no strategic default motive: borrowers with positive home equity. We implement the method using high-frequency administrative data linking income and mortgage default. Our central finding is that only 3 percent of defaults are caused exclusively by negative equity, much less than previously thought; in other words, adverse events are a necessary condition for 97 percent of mortgage defaults. Although this 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.

Peter Ganong
Harris School of Public Policy
University of Chicago
1307 East 60th Street
Chicago, IL 60637
and NBER
ganong@uchicago.edu

Pascal J. Noel
University of Chicago Booth School of Business
5807 South Woodlawn Avenue
Chicago, IL 60637
pascal.noel@chicagobooth.edu

1 Introduction

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

Mortgage defaults soared during the Great Recession, precipitating the worst financial crisis since the Great Depression. As Ben Bernanke explained in the speech quoted above, a key challenge facing policymakers addressing this crisis was understanding *why* defaults were rising. On the one hand, Bernanke notes that defaults could be triggered by an adverse “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. Of course, 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 for lenders and policymakers to develop an appropriate response.

Although Bernanke went on to describe the challenge of distinguishing between life events and negative equity as “novel”, housing economists recognized it as just the latest chapter in a longstanding debate between two theories of mortgage default. The first theory, dating back to Riddiough (1991), is that mortgage default is triggered by a life event that reduces borrower cash flows. The second theory, dating back to Foster and Van Order (1984), is that default is caused by negative equity. According to this theory, borrowers treat their homes like a financial asset. They default when they are so far underwater that the option to sell the home in the future is worth less than their mortgage obligations, irrespective of any life event they may have suffered. 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” or “ruthless” default. It is also possible that both life events and negative equity are necessary for default (Foote, Gerardi and Willen, 2008).

The goal of this paper is to help distinguish between these theories by quantifying the share of defaults that fall into two groups: those that can be causally attributed to adverse life events, and those caused exclusively by negative equity, which are sometimes called “strategic”. This topic has been the subject of substantial prior research. Nevertheless, disentangling the contribution of adverse life events from that of negative equity remains one of the “central questions in this literature” according to a recent review article (Foote

and Willen 2018). The question has remained central in part because of two problems: data limitations and a measurement challenge.

First, administrative mortgage servicing data do not contain information on current income or possible triggering life events. Prior research has typically used measures that are dated, such as the household’s payment-to-income ratio at origination, or geographically coarse, such as the regional unemployment rate. As a result, this work may have inadvertently classified borrowers as strategically defaulting when in fact their default was caused by a life event.

Second, even with data on current income, it is unclear what qualifies as an adverse 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, such as at least 10 percent? Can the drop in income be short-lived, or must it last a specific amount of time, such as at least three months?

We overcome these two challenges with new data and a new empirical method. First, to overcome the data limitation, we link mortgage servicing records with administrative bank account records, both from the JPMorgan Chase Institute (JPMCI). Bank account income and balances are rich (albeit noisy) measures of a household’s financial circumstances. Using mortgage servicing data for the same households, we measure home equity and mortgage default. The linked dataset has 2.9 million borrowers, which is orders of magnitude larger than previously-used data on a household’s financial situation around default.

Second, to overcome the measurement challenge, we use the evolution of income of *above water defaulters* as an empirical benchmark of a default motivated by a life event. Because this group has positive home equity, they should not have any strategic motive to default.¹ Instead, it is logical to assume that their defaults must be driven by life events. Consistent with this view, we show that their income declines sharply beginning six months before default. This decline therefore provides a benchmark for what it looks like when we can be confident that a default is caused by a life event. We can then compare the evolution of income for underwater borrowers to this benchmark.

Our central empirical finding is that the decline in income leading up to default by

¹It may initially be surprising that any borrowers with positive equity ever default, since one might expect they would sell their home or tap into their home equity to avoid missing payments. However, matching frictions make it difficult to sell quickly (Low 2018), and institutional frictions make it difficult for borrowers who are unemployed or liquidity-constrained to quickly access illiquid housing wealth (Boar, Gorea and Midrigan 2017; DeFusco and Mondragon 2018). As a result, both mortgage default and foreclosure are quite common for above water borrowers. Our conclusions are not sensitive to how we define above water default. We define above water default as three missed payments and a loan-to-value (LTV) ratio of less than 100 percent in the baseline specification. However, we allow for alternative missed payment thresholds and LTV thresholds as low as 60 percent in the robustness analysis.

underwater defaulters—whose reasons for default are not known a priori—is visually and statistically indistinguishable from that of above water defaulters—whose defaults must be caused by a life event. The decline in bank account balances is likewise indistinguishable between defaulters with positive and negative equity. This evidence that underwater defaulters experience the same financial distress as above water defaulters is qualitatively consistent with a central role for life events in explaining default.

What causal statement, if any, can be made from this evidence? We make precise the question of interest using potential outcomes (Rubin 1974). We assume that either a life event or negative equity is necessary for default. We label a default as “strategic” when negative equity is a necessary and sufficient condition for the default; “cash-flow” when a life event is necessary and sufficient; and finally, “double-trigger” when both negative equity and a life event are necessary.² We argue that under plausible assumptions, comparing the paths of income for above and underwater borrowers can separate the strategic defaults (where life events are irrelevant) from the cash-flow and double-trigger defaults (where life events are necessary). Yamamoto (2012) calls this type of empirical exercise, which seeks to isolate the fraction of an outcome for which treatment is necessary, *causal attribution*. Measuring what fraction of an outcome is attributable to treatment is a common estimand in economics research.³ This estimand is hard to identify in our setting because we do not observe life events directly.

We propose a novel and transparent method for causal attribution of mortgage default with two ingredients. The first ingredient is a noisy measure of treatment. We assume the noise is orthogonal to the treatment and the outcome, as in the classical errors-in-variables problem (Wooldridge 2010). We use the change in bank account income as a noisy measure of life events. The second ingredient, which is novel, is a group that we assume has been treated with certainty. This is the above water defaulters, who we assume have experienced an adverse life event.

We combine these two ingredients using Bayes rule. The standard approach to causal inference puts the outcome on the left-hand side and treatment on the right-hand side of a regression equation. However, this approach suffers from attenuation bias when treatment is measured with error. In our application, if adverse life events are imperfectly observed at the individual level, simply regressing default on a noisy measure of such events will lead

²The label of “strategic” default as one meant to maximize a borrower’s financial wealth, irrespective of any adverse life event, goes back to Riddiough and Wyatt (1994). This original meaning is commonly used in recent literature (see e.g. Bhutta, Dokko and Shan 2017), and our formal definition follows in this tradition. However, we note that there are other potential definitions of strategic default. We discuss how our results relate to several alternative definitions in Section 5.

³See Section 2 for examples.

researchers to underestimate their importance in driving defaults. We therefore use Bayes rule to move *treatment* to the *left-hand* side, where noise will result in larger standard errors but not attenuation bias. In contrast to the standard approach, which would require studying the outcome conditioning on (noisily measured) treatment, we instead study treatment conditioning on the (precisely measured) outcome.

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

Within the potential outcomes framework, our empirical estimates imply that only 3 percent of underwater defaults are strategic. In other words, we find that adverse life events are a necessary condition for 97 percent of mortgage defaults. Although Bernanke and many others (see e.g. Roubini 2008) were worried about borrowers walking away from their homes solely due to negative equity, our point estimate shows little evidence of this type of default. This estimate has good precision. We calculate that at most 11 percent of underwater defaults in our sample are strategic using a 95 percent confidence interval.

This empirical conclusion holds for different quantiles of income, time periods, and loan-to-value (LTV) cutoffs. First, the result does not depend on our choice of the mean as a summary statistic. It holds across the income distribution. Second, we see almost no strategic default in every year available in our data, including 2010, which is when the national default rate reached its peak. Third, this finding does not depend on choosing an LTV cutoff at exactly 100 percent, as the same result holds for cutoffs ranging from 60 percent to 160 percent.

This finding of almost no strategic default is potentially surprising because some prior empirical work estimates that between 30 and 70 percent of Great Recession defaults are strategic (Gerardi et al. 2018; Guiso, Sapienza and Zingales 2013; Bhutta, Dokko and Shan 2017). Why is our estimate so much lower? One candidate explanation is that there is something unusual about the administrative data we use. A second candidate explanation is that prior research has used a range of definitions for what constitutes a strategic default. A third candidate explanation is that error in measuring life events may have led to attenuation bias in prior research designs such that they underestimated the importance of such events.⁴

⁴The literature analyzing regional unemployment rates and default offers one example of how attenuation bias may arise when measuring the contribution of adverse life events to default. A long literature beginning

To help adjudicate between these candidate explanations, we analyze publicly available data from the Panel Study of Income Dynamics (PSID).

Applying our method to PSID survey data yields similar conclusions about the lack of strategic default; measurement error in life events explains why our estimates differ from prior work. Although the administrative data is key for obtaining precise estimates of the prevalence of strategic default, we show that similar patterns are apparent in the survey data. Furthermore, this exercise allows us to compare our estimates to prior methodologies and prior definitions of strategic default based on the PSID. 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.

Our finding of almost no strategic default may also be surprising because existing structural models predict substantial strategic default by deeply underwater borrowers. We use a benchmark structural model of mortgage default first developed in Campbell and Cocco (2015) to illustrate this point. This type of model is particularly useful for our research question because it is the first to integrate strategic motives based on option-value theory and cash flow motives based on realistic income risk. Even though individual-level data on income around default was not available as a target for the development of this model, we find that the model’s predictions closely match the data for borrowers with LTV less than 120. However, as LTV rises above this threshold, the model predicts that borrowers will default even in the absence of income shocks, which contrasts sharply with our empirical findings.

Despite this divergence between model and data for deeply underwater borrowers, we show that a simple extension can reconcile the two. Specifically, this type of model allows for the possibility that defaulting incurs a utility cost. However, Campbell and Cocco explain that the main difficulty with incorporating this cost is that there has previously been little data to discipline this parameter.⁵ We propose to use our new moments on income losses before default as empirical targets. Intuitively, the reluctance of borrowers to default on a substantially underwater asset in the absence of income shocks is informative about how

with Campbell and Dietrich (1983) finds that regional unemployment has modest predictive power, which has been interpreted as consistent with a large role for strategic default. In the recent crisis, Goodman et al. (2010) used regional unemployment and titled their study: “Negative equity trumps unemployment in predicting mortgage default”. Yet as researchers have acquired better micro data and developed improved simulation methods, some evidence has emerged suggesting that unemployment may 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, Hsu, Matsa and Melzer (2018) show that mortgage default by the unemployed is highly responsive to cash-on-hand.

⁵Indeed, prior work has estimate a wide range of costs, from as low as a 1.5 percent decrease in the constant-equivalent consumption stream to as high as a 70 percent decrease (Kaplan, Mitman and Violante 2017; Hembre 2018; Schelkle 2018; Laufer 2018).

costly they perceive this default to be. We estimate that defaulting must impose a cost equivalent to a 25 percent decrease in the constant-equivalent consumption stream.⁶ Once this cost is incorporated, we find that the model is able to closely match the data. A high default cost thus provides one plausible microfoundation for the behavior we observe empirically.

This approach to reconciling model and data may provide a blueprint for a wide class of macro finance models where borrower default decisions play a central role. For example, models with endogenous borrower default have been used recently to inform questions about macro-prudential regulation, the origins of the 2008 financial crisis, bankruptcy and foreclosure policy, and optimal mortgage security design.⁷ These types of models must take a stand on what triggers borrower default. Our empirical results suggest that realistic models will feature life events such as cash flow shocks to be a necessary condition for most default. We demonstrate that incorporating a large utility cost of defaulting is one specific way to achieve this. More generally, regardless of exactly how the default decision is modeled, in order to match the data, models with endogenous borrower default might seek to target large income drops before default, even for deeply underwater borrowers.

Our results have two types of policy implications. First, they relate to questions about how easy it should be for lenders to foreclose on delinquent borrowers. The penalties for mortgage default can be analyzed as an optimal social insurance problem where the planner seeks to equalize the consumption-smoothing benefits of default with the concern that strategic defaults raise the cost of credit (Dávila 2016). Policymakers have several levers that might affect strategic default, including whether to require judicial review of foreclosures, setting the length of time to foreclose, and allowing recourse to seize non-mortgage debt. But if strategic default is rare because the private cost of mortgage default is high, then stringent public penalties may be unnecessary.⁸

Second, our results on the lack of default driven by negative equity bolster an active literature documenting the superiority of payment reduction to principal forgiveness as a foreclosure prevention tool (Eberly and Krishnamurthy 2014; Ganong and Noel 2018; Scharlemann and Shore 2019). The policy design questions addressed in this literature may again become crucial in response to COVID-19, as April 2020 saw the largest single-month increase in mortgage delinquency ever recorded (Black Knight 2020). The contribution of this paper to

⁶We discuss potential sources of this cost in Section 6.

⁷See e.g., Corbae and Quintin (2015), Mitman (2016), Kaplan, Mitman and Violante (2017), Guren, Krishnamurthy and McQuade (2018), Campbell, Clara and Cocco (2018), Greenwald, Landvoigt and Van Nieuwerburgh (2018), Diamond and Landvoigt (2019), and Garriga and Hedlund (2020).

⁸In this way, our results contribute to an active literature that analyzes the effect of variation in penalties for debt non-repayment on default (Ghent and Kudlyak 2011; Gross et al. 2019; Indarte 2019; O'Malley 2018).

that literature is to study borrowers that the previous literature classified as especially likely to default strategically. Prior models and empirical evidence suggest that strategic default will be most prevalent for deeply underwater borrowers. Yet policy evaluations of principal forgiveness have analyzed borrowers who are moderately underwater or above water. In addition, many papers rely on variation arising from mortgage modifications, which often have stringent eligibility criteria designed to exclude strategic defaulters. We address these limitations using a sample that includes deeply underwater borrowers and borrowers who did not receive modifications.

The paper proceeds as follows. Section 2 describes a novel econometric framework for causal attribution. Section 3 describes the two datasets with income, home equity, and default (administrative bank data and PSID survey data). Section 4 shows that life events are a necessary condition for almost all defaults using the administrative data, which is contrary to theories of strategic default. Section 5 replicates this finding using the PSID. Section 6 explores implications for modeling mortgage default. Section 7 concludes.

2 Econometric Framework

2.1 Standard Approaches to Causal Attribution

In many social science applications, researchers seek to measure “how much of event Y is attributable to binary treatment T^* ?” One precise answer to this question is the change in outcome Y from eliminating treatment T^* , which can be written as

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

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

In this section we describe a new method for causal attribution. We first describe challenges with existing approaches. Second, we describe a method that addresses these challenges. Third, we describe how the method can be applied to study mortgage default. Finally, we discuss some of the limitations of the approach, and compare and contrast the method to other estimation strategies for causal inference (e.g., difference-in-difference, instrumental variables, reverse regression).

The causal attribution parameter, which we define in equation (1), is of interest in many economic applications. Recent examples that estimate α include: “how much of the correlation in wealth between parents and their children is due to nature versus nurture?” (Black et al. 2019; Fagereng, Mogstad and Ronning 2018), “what fraction of bankruptcies are caused by medical expenses?” (Gross and Notowidigdo 2011; Himmelstein et al. 2005; Dobkin et al. 2018), “what fraction of house price movements during the Great Recession is explained by a shift in beliefs?” (Kaplan, Mitman and Violante 2017), and “what share of layoffs in the Great Recession can be accounted for by credit supply contractions?” (Chodorow-Reich 2014).

There are two standard methods for estimating α : counterfactuals in structural models and “back-of-the-envelope” calculations. The first method, as in the Kaplan, Mitman and Violante (2017) example in the prior paragraph, is to specify an economic environment and utility functions, estimate preferences and other parameters, and evaluate a counterfactual in a structural model using the estimated parameters. The second method relies on the observation that, because $Y = Y_1T^* + Y_0(1 - T^*)$, α can be rewritten as proportional to the product of the average treatment effect and the probability of treatment:

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

This method is used in the other examples in the prior paragraph. In these other examples, researchers estimate a local average treatment effect ($E(Y_1 - Y_0|complier)$), assume that it extrapolates to the population average treatment effect on the treated ($E(Y_1 - Y_0|T^* = 1)$), multiply by an estimate of the probability of treatment ($P(\hat{T}^*)$) and divide by an estimate of the average level of the outcome ($E(\hat{Y})$). This is sometimes called a “back-of-the-envelope” calculation. Applied to analyzing the fraction of mortgage default attributed to life events, this method would require an estimate of the average treatment effect of life events on the probability of default (for example, an increase of 10 percentage points), an estimate of the probability of life events (e.g, 30 percent) and a population default rate (e.g., 6 percent). Multiplying the first two terms and dividing by the third, a researcher could hypothetically conclude that 50 percent of default was attributable to life events.

However, measurement error can bias estimates of α both because of attenuation bias in the estimate of treatment effects and because the probability of treatment is unknown. This challenge emerges in many social science questions, because we often do not see the explanatory variables we want and instead we see noisy proxies. For example, Allcott, Lockwood and Taubinsky (2019) seek to quantify “how much of soda consumption is caused by self-control problems?” They measure self-control over soda consumption using a survey

where responses can take one of four categorical values. Self-control would be measured with error if people are not perfectly aware of their own self-control problems or if the underlying self-control parameter is continuous.

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

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

2.2 A New Method for Causal Attribution

We propose an alternative method to estimate α . The standard “back-of-the-envelope” approach to causal inference puts the outcome on the left-hand side and treatment on the right-hand side of a regression equation. However, this approach suffers from attenuation bias when treatment is measured with error. We therefore use Bayes rule to move *treatment* to the *left-hand* side, where noise will result in larger standard errors but not attenuation bias.

Our method requires two ingredients. The first ingredient is a noisy measure of treatment. We assume the noise is orthogonal to the treatment and the outcome, as in the classical errors-in-variables problem (Wooldridge 2010). We use the change in bank account income as a noisy measure of life events in our application. We describe why this measure is noisy in the data description in Section 3. The second ingredient, which is novel, is a group that we assume has been treated with certainty. In our application, this is the above water defaulters, whom we assume are treated (e.g., they must have experienced a life event).

Environment and assumptions

We assume there exists a population distribution (T^*, G, Y, T) where T^* is treatment, G is group status, Y is the outcome, and T is a candidate noisy measure of T^* . The first three variables are binary. In our mortgage default application, $T^* = 1$ indicates a life event, $G = 1$ is negative home equity, $Y = 1$ is mortgage default, and T is the change in bank account income relative to one year ago. Assume potential outcome function $Y(T^*, G)$ (sometimes denoted as Y_{T^*G} for brevity), so each individual has four potential outcomes: $Y(0, 0)$, $Y(1, 0)$, $Y(0, 1)$, and $Y(1, 1)$.⁹ Additionally, assume there exists a potential outcome function $T(T^*, G, Y)$. The econometrician observes random draws from (G, Y, T) but T^* is unobserved.

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

In our application, this assumption says that the outcome of mortgage default requires either negative equity or a life event. This implies that a mortgage default *without* negative equity *must* have a life event ($P(T^* = 1 | Y = 1, G = 0) = 1$). The intuition is that a home with positive equity has financial value to the borrower, so defaulting is not in their long-term financial interest. Instead, it must reflect the impact of an adverse shock which reduces the borrower’s ability to afford their mortgage payment. This is the novel assumption that distinguishes this paper from prior work on identification in the presence of classical errors-in-variables (CEV). This assumption is similar in spirit to papers in the nature-versus-nurture literature that study adoptees in order to isolate the role of nurture by shutting down the nature channel (Kleven, Landais and Sørensen 2020).

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

The idea that above water default ever occurs may be surprising to some readers. After

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

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

all, if houses were a completely liquid asset, then above water default would never occur because a borrower could borrow against their home to cover the mortgage payment.

Yet in practice, there are substantial frictions to accessing home equity for borrowers in financial distress (Boar, Gorea and Midrigan 2017; DeFusco and Mondragon 2018). Second lien underwriting and refinancing require a good credit history and a documented “source of repayment” (Office of the Comptroller of the Currency 2005), which usually means proof of income or proof of substantial liquid assets (Fannie Mae 2011). In the Fannie Mae underwriting guide, unemployment insurance is not an acceptable source of income. An unemployed homeowner who needs a loan to cover her current mortgage payments would not meet the prevailing underwriting standard during our sample period.

Borrowers may also choose to sell their home, but there are frictions in this process as well. Low (2018) shows that above water default can be rationalized by a quantitative model with matching frictions in the home sale market that make the time-to-sell and resale value uncertain. Consistent with Low’s assumption about such frictions, Guren (2018) documents that fewer than half of listed homes were sold within three months.

These two frictions mean that above water default is quite common. In online Appendix Figure 1, we corroborate Low’s finding that above water default is ubiquitous. The figure shows that even at the peak of the housing crisis, 40 percent of defaults were by above water borrowers.

Above water default is costly for borrowers, both because of its immediate credit score impact and because of the risk of eventual foreclosure. The credit score decline from falling behind by three months on a mortgage is more than 80 percent as large as the decline from foreclosure and almost 60 percent as large as the decline from bankruptcy (Christie 2010). We calculate that 55 percent of above water borrowers who fall behind by three months have a foreclosure initiation within one year in online Appendix Table 1. This is slightly below the foreclosure initiation rate of 57 percent for underwater borrowers.¹¹

Assumption 1 serves as a substitute for validation data in enabling us to identify the relationship between observable T and unobservable T^* . In a standard validation data setup, the researcher *knows* treatment with certainty for a subset of the data (Ridder and Moffitt 2007). In contrast, Assumption 1 says that there is a group where we can *assume*

¹¹A priori, it is not obvious whether lenders should foreclose more often on above water or underwater homes whose mortgages are in default. One reason why underwater foreclosure should occur more often is that the lender may hope that above water defaulters will sell their homes, enabling the lender to avoid substantial administrative costs. On the other hand, the return to the lender from foreclosing on an above water home is higher; it is possible that even after accounting for administrative costs they will be able to recoup the full balance outstanding on the loan. In contrast, a lender who forecloses on an underwater home knows they will not be able to recoup the full balance, and so may have an incentive to wait and hope the borrower repays the original loan instead.

treatment with certainty for observations with $Y = 1$ and $G = 0$ using an economically sensible restriction on the potential outcome function.

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

This assumption says treatment is exogenous with respect to the potential outcome $Y(T^*, G)$ conditional on group G and is standard in the literature on measurement error. In our mortgage default application, this assumption rules out omitted variable bias where, after conditioning on home equity, there is a third factor which causes both a life event and also causes mortgage default. This assumption is also made in the CEV framework in Wooldridge (2010).

This assumption allows for three types of heterogeneity that are important in the mortgage default context. First, it allows for the possibility that underwater borrowers are more likely to have adverse 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 2017.¹² Second, it allows for heterogeneity in the causal impact of a life event on default, consistent with findings in Gerardi et al. (2018) that underwater borrowers are more sensitive to income shocks than above water borrowers ($E(Y_{11} - Y_{01}) > E(Y_{10} - Y_{00})$). Third, it does *not* require that G is exogenous with respect to potential outcomes; instead, the role of G is similar to a covariate that we condition on in the analysis.

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

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

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

An implication of Assumption 3 is that when a life event *does* occur, above and underwater borrowers have the same average decline in income. This assumption allows us to use T

¹²This is also consistent with evidence in Bernstein (2019) and Gopalan et al. (2019), who find that borrowers with negative equity are more likely to suffer income declines because of constrained mobility and financial distress.

of above water defaulters, who always have $T^* = 1$ by Assumption 1, to learn about $P(T^*)$ for underwater defaulters.

Assumption 3 would fail if income for above water defaulters who experienced a life event fell more or less than income for underwater defaulters who experienced such an event. The first case might occur if above water borrowers with smaller income drops used alternative forms of credit to avoid default, while underwater borrowers did not. The second case might occur 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. These scenarios have a testable prediction: 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. We evaluate this prediction in Section 4.3; the data do not support this prediction.

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

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

We use these four assumptions to interpret the default motivations of underwater borrowers. Specifically, we ask what fraction of underwater mortgage default can be causally attributed to life events? This corresponds to estimating α for the $G = 1$ group in the model, which we call $\alpha_{G=1}$

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

$$\alpha_{G=1} \equiv \frac{E(Y|G=1) - E(Y_{01}|G=1)}{E(Y|G=1)} = \frac{E(T|Y=1, G=1) - E(T|G=1)}{E(T|Y=1, G=0) - E(T|G=1)}. \quad (2)$$

Proof: See Appendix C.

The key step in the proof relies on Bayes rule. The standard approach to causal attribution puts the outcome Y on the left-hand-side and treatment T^* on the right-hand side of a regression equation. However, this approach suffers from attenuation bias when treatment is measured with error (i.e., we observe T instead of T^*). We therefore use Bayes rule to move the measure of treatment T to the *left-hand* side, where noise will result in larger standard errors but not attenuation bias. Hence, equation (2) relies on comparing this noisy measure of treatment across groups. Below, we provide additional intuition for the formula in the context of mortgage default.

2.3 Application to Mortgage Default

2.3.1 Formula and intuition

In the mortgage default application, our method is easy to use and interpret. We can write equation (2) as:

$$\alpha_{\text{life event underwater}}^{\text{life event}} = \frac{\overbrace{E(\Delta Inc^{\text{UnderwaterDefaulter}})}^{\text{group of interest}} - E(\Delta Inc^{\text{AllUnderwater}})}{\underbrace{E(\Delta Inc^{\text{AbovewaterDefaulter}})}_{\text{benchmark of 1: all defaults caused by life events}} - \underbrace{E(\Delta Inc^{\text{AllUnderwater}})}_{\text{benchmark of 0: no defaults caused by life events}}} \quad (3)$$

Going forward, we refer to $\alpha_{\text{life event underwater}}^{\text{life event}}$ as α for brevity.

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

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

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

Because income is “irrelevant” for this decision, it is natural to expect defaulters motivated by negative equity to have the same change in income as non-defaulters. Thus, if we observe the same income drop for underwater defaulters and nondefaulters, i.e. $E(\Delta Inc^{\text{UnderwaterDefaulter}}) = E(\Delta Inc^{\text{AllUnderwater}})$, then our assumptions imply that adverse life events play no role in the default decision ($\alpha = 0$). Finally, if $\Delta Inc^{\text{UnderwaterDefaulter}}$ is in between the two benchmarks, then the share of defaults causally attributed to life events is between 0 and 1.

2.3.2 Potential Outcomes Interpretation

Our framework also allows us to separate default behavior in terms of four potential outcome types which correspond to specific economic mechanisms in the context of mortgage default. Recall that mortgage default Y is a function of two binary variables, life event treatment T^* and equity status G , i.e $Y_{\text{life event, negative equity}}$. We separate defaulters into four types based on possible combinations of these potential outcomes.¹³

These types have natural labels in terms of the prior literature, which we show in Table 1.

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

The estimand α has a clear interpretation relative to these potential outcome types. As defined in equation (2), α measures the fraction with potential outcome $Y_{01} = 0$ (i.e., the fraction of underwater defaults that would be eliminated in the absence of life events). It therefore quantifies the combined share of cash-flow and double-trigger defaults (CF + DT) among underwater defaulters. These are defaults where life events are a *necessary* condition,

¹³The fact that there are four potential outcome types arises from Assumptions 1 and 4. In principle, with four binary potential outcomes for each household $\{Y_{00}, Y_{01}, Y_{10}, Y_{11}\}$, there are $2^4 = 16$ possible combinations. We narrow the set of potential outcomes with two assumptions. First, Assumption 1 is that default requires either a life event or negative equity, so $Y_{00} = 0$ for all borrowers, which leaves $2^3 = 8$ possible combinations for the remaining three binary potential outcomes. Second, Assumption 4 (monotonicity) rules out any combination where Y_{01} or Y_{10} are 1 (i.e. the borrower would default with only a life event or with only negative equity) but $Y_{11} = 0$ (i.e. the borrower wouldn’t default with both a life event and negative equity). This leaves five potential outcome types. There is one who never defaults where all potential outcomes are zero regardless of life events or negative equity. This leaves four potential outcome types among defaulters.

or, in the language of Yamamoto (2012), defaults that can be “causally attributed” to life events.

The balance, $1 - \alpha$, quantifies the share of underwater defaults for which negative equity is a *sufficient* condition. These are the borrowers with potential outcome $Y_{01} = 1$. Although this quantity technically combines both the strategic and the sensitive types (ST + SE), we will refer to $1 - \alpha$ as the share of strategic defaults in most of the subsequent analysis. This is because we are estimating this object specifically for underwater borrowers. Since these borrowers have negative equity, and the sensitive type would default solely due to this negative equity ($Y_{01} = 1$), life events are effectively irrelevant to their default decision. However, the presence of this type means that $1 - \alpha$ is an upper bound on the subset that might be considered canonical strategic defaulters, where negative equity is both necessary and sufficient. Hence, if anything, the share accounted for by this canonical strategic group may be even smaller than $1 - \alpha$.

While our estimand $\alpha_{\text{underwater}}^{\text{life event}}$ from equation (3) captures the fraction of defaults for which life events are a necessary condition, a related estimand captures the fraction of defaults for which negative equity is a necessary condition (call this $\alpha_{\text{underwater}}^{\text{negative equity}}$). While $\alpha_{\text{underwater}}^{\text{life event}}$ measures the fraction with $Y_{01} = 0$ (CF + DT), $\alpha_{\text{underwater}}^{\text{negative equity}}$ would measure the fraction with $Y_{10} = 0$ (ST + DT). The $\alpha_{\text{underwater}}^{\text{negative equity}}$ estimand therefore measures the combined fraction of strategic and double-trigger defaults, as depicted in Table 1.

It is useful to emphasize that while both $\alpha_{\text{underwater}}^{\text{life event}} \in [0, 1]$ and $\alpha_{\text{underwater}}^{\text{negative equity}} \in [0, 1]$, their sum could be greater than 1. This is because double-trigger defaults are captured by both estimands. As we discuss above in the context of Assumption 2, we do not impose additive separability between life events and negative equity. Rather, we allow for the potential that there is a group that would not default in the presence of only one trigger, but would in the presence of both. This is a feature of any causal attribution exercise that seeks to decompose the role of two different channels where interaction effects are allowed. For example, returning to the classic nature versus nurture debate, many researchers now estimate specifications that allow for interactions between the effects of nature and nurture (e.g. Black et al. 2019).

Although we think $\alpha_{\text{underwater}}^{\text{negative equity}}$ is also an interesting object, it is not identified by our methodology and we do not attempt to measure it in this paper. Instead, our estimand focuses on separating the strategic defaults from the cash-flow and double-trigger defaults. This separation is helpful for understanding borrower behavior and for designing policies to address mortgage default. From a policy perspective, if default is strategic and therefore driven exclusively by negative equity, then reducing that negative equity by reducing mortgage principal may be necessary to prevent default. However, principal reduction is expensive for lenders (and, if the policy is subsidized, for taxpayers). On the other hand, if

instead defaults are either cash-flow or double-trigger, then eliminating the cash-flow motive is sufficient for preventing default. This may be accomplished much more cheaply by offering, for example, forbearance or mortgage term extensions, which provide immediately liquidity while leaving the principal balance unchanged (Ganong and Noel 2018, Campbell, Clara and Cocco 2018). It is therefore crucial to know what fraction of defaults are driven exclusively by negative equity and what fraction are influenced by life events such as cash-flow shocks.

2.4 Limitations and Comparison to Prior Literature

The benefit of our approach is that it offers unbiased causal attribution estimates when treatment is measured with error (and so conventional methods would suffer from attenuation bias). However, the approach has two types of limitations worth noting. First, it identifies fewer parameters of interest than the standard approach to measurement error. Although it identifies the product of the treatment effect and the probability of treatment, it does not distinguish between the two. For example, a finding of $\alpha = 1$ in our application could be consistent with either infrequent adverse life events that have a large impact on default or frequent ones that have a small impact on default. Because the main problem is that life events themselves are difficult to measure, we cannot distinguish between these two. To identify the treatment effect, in addition to Assumptions 1-4, we would also need to know the probability of treatment.

The second limitation is that the method imposes additional restrictions on the data-generating process which are not needed for the standard approach. First, it requires that the outcome variable be binary and measured without error. These requirements are essential because we use a combination of group status and the outcome to infer treatment. Second, it requires that the treatment can be modeled as binary. The appropriateness of this assumption is context-specific. In the mortgage default context, Guiso, Sapienza and Zingales (2013), Bhutta, Dokko and Shan (2017) and Gerardi et al. (2018) have modeled treatment as binary. In addition, we show that it is possible to offer a similar interpretation of the data without the binary treatment assumption using a structural model of mortgage default in Section 6.

Our approach has some parallels with reverse regression, with the difference-in-differences (DiD) design, and with instrumental variables (IV). It is similar to reverse regression in that we correct for attenuation bias in a regression by moving the noisy variable from the right-hand side to the left-hand side of the regression equation.¹⁴ However, reverse regression is

¹⁴The classic example of reverse regression is a researcher who seeks to measure the whether an employer is discriminating against their female employees in setting wages. The ideal test is to regress wages on

most commonly used to reject or fail-to-reject a specific null hypothesis (e.g., “is there wage discrimination?”), while our approach provides a quantitative, causal interpretation of the sample moments.

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

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

3 Data

One of the key impediments to understanding why borrowers default is a lack of data. Mortgage servicing datasets do not include information that captures a borrower’s financial circumstances at the time of default. A review article by Vandell (1995) says that to make progress on mortgage default, researchers should “develop a microbehavioral mortgage payment data” and then “gather detailed information whenever termination occurs.”

Progress on this front has been slow. In a review article published over two decades later, Foote and Willen (2018) call for constructing almost exactly the same dataset.

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

3.1 Administrative Data

We address this challenge using a novel administrative dataset from Chase that links checking account records and mortgage servicing records. These records are linked and then de-identified by Chase.¹⁵ This linkage is possible because Chase is both a consumer bank that offers checking accounts and a mortgage servicer.

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

A second key variable—the loan-to-value (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.¹⁶ Estimated home value is constructed using the standard procedure of inflating purchase price by a local measure of house price changes from CoreLogic. How accurate is this procedure? Giacoletti (2017) finds that the standard deviation (σ) of the error in estimated home values (as compared to actual sale prices) is 18.7 percent. Although in our primary analysis we define positive home equity as measured $LTV < 100$ percent, in our robustness analysis we also report results for a subsample with measured $LTV < 60$ percent. This subsample is *very* likely to have positive home equity; using Giacoletti’s estimate we calculate that it would require an error of three standard deviations (3σ) for such a borrower to actually have negative home equity.

We study borrowers who have cumulatively fallen behind on their mortgage by three monthly payments in most of our analysis. This is also known as 90 days past due. This is a common threshold for a mortgage to be considered in default (Foote and Willen 2018; Bhutta, Dokko and Shan 2017). However, we examine other thresholds for default in our robustness analysis. Although we do observe whether the lender has initiated foreclosure proceedings, Bhutta, Dokko and Shan (2017) argues that this is endogenous to the *lender* and therefore default is a better measure of *borrower* behavior. Nevertheless, as a robustness check we implement our test for the subset of borrowers whose lenders initiate a foreclosure proceeding. We also observe whether the borrower self-identifies as an investor and whether

¹⁵See Farrell et al. (2017) and Farrell, Bhagat and Zhao (2018) for JPMCI research using this linked dataset.

¹⁶Nineteen percent of Chase-serviced first liens had second liens in 2011, which is similar to 15 percent for a benchmark sample of first liens linked to credit bureau data called Credit Risk Insight Servicing McDash.

this home will be the borrower’s primary residence at mortgage origination.

Our primary analysis uses first lien mortgages serviced by the bank between October 2012 and August 2015 and household income measured using checking accounts. The linked dataset has 2.934 million mortgages. We analyze defaults with an open checking account from one year before default through the date of default. This analysis sample contains 29,433 defaults (see online Appendix B for details on sample construction). We also analyze the evolution of income for the universe underwater borrowers, both defaulters and non-defaulters (408,000 borrowers). Unfortunately, data on income are unavailable prior to October 2012.

We also use a second sample with data on checking account balances, which (unlike income) are available between January 2007 and August 2015. This dataset has 5.032 million borrowers. We apply the same sample screens to the balance data as described above. This produces an analysis sample of 138,345 loans in default. The analysis sample is a larger proportion of the total sample because the mortgage default rate was higher from 2007 to 2012. We also analyze the evolution of balances for the universe of underwater borrowers, both defaulters and non-defaulters (1.035 million borrowers).

Mortgage borrowers at Chase have similar characteristics to mortgagors in other, more widely-used datasets. We compare the borrowers in our data to the McDash and Credit Risk Insight Servicing McDash (CRISM) data. McDash is a dataset of origination and servicing records which covers about 70 percent of outstanding mortgage balances during our sample period. CRISM is a subset of the McDash data that can be linked to individual credit bureau records (and therefore can be used to calculate LTV ratios), which covers about 50 percent of outstanding mortgage balances during this time period. Online Appendix Table 1 shows that the Chase data are similar to McDash and CRISM in terms of the default rate, the share of borrowers who are investors, the share of all borrowers who are underwater, and the share of defaulters who are underwater.

3.2 Survey Data

To check that our results are not unique to the Chase sample, we also conduct parallel analysis using the best available public use data on income and mortgage default, which is survey data from the Panel Study of Income Dynamics (PSID). The PSID records pre-tax income y and consumption c in the calendar year prior to the survey. A supplement administered in 2009, 2011, and 2013 records housing costs (mortgage, property taxes, and insurance) m , home equity, and default (measured as 60 days past due) at the time of the survey. We follow the sampling choices from Gerardi et al. (2018): we drop households with

LTV > 250 percent, and we require that the head of household is in the labor force and between the ages of 24-65.

The PSID has two benefits relative to the JPMCI data. First, it captures a sample from all bank accounts and mortgage servicers (rather than the universe from people who have checking accounts at, and a mortgage serviced by, one bank). Second, it captures borrowers' *perceived* LTV, which is the decision-relevant measure of LTV from the borrower's perspective, circumventing possible concerns about measurement error in LTV described above. However, income is observed only once every two years and there are only 268 households that default on their mortgages, which is two orders of magnitude smaller than the sample size of defaulters in the JPMCI data.

4 Main Results: Estimates of Defaults Causally Attributed to Life Events from Administrative Data

Our research design compares the evolution of mean income around default by home equity. We find that income declines are indistinguishable for above and underwater defaulters. Using the econometric framework in Section 2, this empirical result implies that almost all defaults are causally attributable to adverse life events.

4.1 Research Design

Our method in equation (3) compares underwater defaulters' income to two benchmarks. In this section we discuss each of the two benchmarks and how the patterns they show relate to the assumptions underlying our approach. Then in Section 4.2 we put them together to calculate an estimate of the share of defaults causally attributable to life events.

Figure 1a shows the evolution of income in the twelve months prior to mortgage default. The x-axis is months since three missed payments. The figure is similar to an event study in that it shows monthly data relative to an event. It differs in that we focus specifically on income data *prior* to default. This choice is motivated by this paper's focus on the *causes* of the event, whereas the traditional event study is usually interested in the *consequences* of the event. The dependent variable is the change in monthly income relative to average income one year before default.

This figure is useful for evaluating some of the identifying assumptions from Section 2. The timeseries of income is consistent with Assumption 1, which says that above water defaults are caused by life events. Income falls sharply in the two quarters prior to default.

Online Appendix Figure 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.¹⁷

Our econometric method assumes that income is an equally good way to measure adverse life events for above water and underwater borrowers (Assumption 3). This identifying assumption has some testable implications. Intuitively, the assumption would fail if these two groups differ in terms of their economic fundamentals in the year prior to default or if one group differentially manipulates their income. To gauge the magnitude of the former concern, we compare the trends in income of above water versus underwater borrowers more than six months prior to default, which is when the life event appears to occur for above water borrowers. This is similar to the test for parallel pre-trends in an event study research design. Figure 1a shows that income trends similarly during this time period. We interpret this to indicate that above water defaulters are a good comparison group for underwater defaulters in terms of economic fundamentals.

Another possible threat to Assumption 3 would arise if underwater borrowers hide some of their income from the bank that services their mortgage.¹⁸ We address this concern in two ways. First, we investigate whether defaulters redirect their Social Security income away from their Chase bank account, and we find no evidence of such behavior (see online Appendix Figure 3).¹⁹ Second, we conduct a parallel analysis in Section 5 using a dataset where borrowers have no motive to manipulate their income (the PSID).

Finally, we emphasize that Assumption 3 does *not* depend on assuming that above water and underwater borrowers have the same degree of financial vulnerability. Table 2 shows that underwater borrowers have slightly higher income levels, bank account balances, and mortgage payment due than above water borrowers. The key assumption is that income declines by the same amount in percent terms *conditional on a life event*.

Figure 1a also shows the evolution of income for all underwater borrowers (both defaulters and non-defaulters). This series provides the benchmark for what income would look like if all defaults were strategic. We construct this series by re-weighting average income by

¹⁷However, it would not be a problem for our methodology if the defaults in our sample were triggered by changes in payments. Payment changes are a non-negative-equity channel that could cause default and are thus considered a “life event” in the context of our model. Furthermore, payment changes directly affect borrower ability to pay.

¹⁸This concern might arise because of two different types of borrower misperception. First, if a borrower incorrectly believes that assets at the mortgage-servicing bank are more likely to be seized than assets at another bank. Second, manipulation could also arise if the borrower incorrectly believes that lower bank account income will lead to a more generous mortgage modification. However, such beliefs are inaccurate. In fact, seizure can only occur after a foreclosure is completed and a deficiency judgment has been rendered. In addition, the bank’s publicly-available mortgage assistance form asks for paystubs and income as reported to tax authorities; it does not ask for bank account income.

¹⁹Social Security income is generally stable, and so if we observed a decline it would be natural to conclude this was due to manipulation rather than an actual lapse in benefits.

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

$$Income^{AllUnderwater} = \sum_s Income_s^{AllUnderwater} w_s. \quad (4)$$

To capture average income of all underwater in months *prior* to default, we compute $Income_t^{AllUnderwater} = \sum_s Income_{s+t}^{AllUnderwater} w_s$ where t is the number of months until default for $t \in \{-12, -11, \dots, 0\}$. Figure 1a shows that this series increases by 0.8 percent over the year, perhaps reflecting inflation. The series also exhibits a modest S-shape—rising through month -8, falling through month -4, and then rising again—which is driven by correlated seasonal patterns in income and mortgage default (see online Appendix Figure 4 for details).

4.2 Central Estimate

Our central empirical result—that the evolution of income is indistinguishable for underwater and above water defaulters—is shown in Figure 1a. The figure shows that income drops for underwater defaulters just as much as for above water defaulters. To provide quantitative estimates on defaulters, we analyze data for a three month “pre-period” well before default ($t = \{-12, -11, -10\}$) and around the time of default ($t = \{-2, -1, 0\}$). We regress

$$\frac{Income_t}{Income_{pre}} = \lambda + \kappa(LTV > 100) + \underbrace{\gamma \mathbf{1}(t = -2, -1, 0)}_{\text{abovewater drop at default}} + \underbrace{\beta \mathbf{1}(t = -2, -1, 0) \times LTV > 100}_{\text{difference for underwater}} + \varepsilon \quad (5)$$

where $Income_{pre}$ is average income in the pre period, computed separately for above and underwater defaulters. Table 3 column (1) shows that the above water income drop $\hat{\gamma}$ is 0.203. This means that monthly income falls by 20 percent in the month of mortgage default and the two months prior. Table 3 also shows that $\hat{\beta}$ is 0.006 with a standard error of 0.007. Thus, we are unable to reject $\beta = 0$ —the hypothesis that the drop for underwater defaulters is the same as the drop for above water defaulters—using a 95 percent confidence interval.

Applying the framework from Section 2 to our regression estimates, we find that 97 percent of defaults are causally attributable to life events. Equation (3) requires three inputs to estimate α : 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 $\hat{\varphi}$). We estimate that

$$\hat{\alpha} = \frac{(\hat{\gamma} + \hat{\beta}) - \hat{\varphi}}{\hat{\gamma} - \hat{\varphi}} = \frac{-0.203 + 0.006 - 0.028}{-0.203 - 0.028} = 97\% \quad (6)$$

with a standard error, computed using the delta method, of 4 percent, as shown in Table 4a. The lower bound of the 95 percent confidence interval on the share of defaults causally attributable to adverse life events ($\hat{\alpha}$) is 89 percent. The interval also includes $\hat{\alpha} = 100$ percent.

We use the potential outcomes model in Table 1 to interpret this point estimate. Ninety-seven percent of underwater defaults have potential outcome types that we categorize as “cash flow” or “double-trigger”, while at most ($1 - \hat{\alpha} = 1 - 0.97 =$) 3 percent of defaults are “strategic”.

4.3 Robustness

Our finding of very little strategic default is consistent across several specifications shown in Table 4a and 4b.

Distribution of Income Changes For example, the finding does not hinge on the choice of a mean as a summary statistic; in fact, it holds across the entire distribution of the change in income. To demonstrate this, the histogram in Figure 1b shows that the distribution of the change in monthly income is visually indistinguishable for above and underwater borrowers.²⁰ Table 3 similarly shows in columns (2)-(4) that the 25th, 50th, and 75th percentile of the income change distribution are statistically indistinguishable for above and underwater borrowers. Our quantile-based estimates of α shown in Table 4a are consistent with very little strategic default. In Section 2.2, we discussed the possibility that the similarity of the mean income drop between groups might be masking distributional heterogeneity where underwater defaulters experience different drops in income. However, the distribution in Figure 1b shows no evidence of such heterogeneity.

One notable feature of Figure 1b 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

²⁰Online Appendix Figure 5 shows the same lesson using the cumulative distribution function.

to cover an unexpected expense.²¹

Definition of Default Our results are not sensitive to the definition of default. Although theoretical models of default typically treat it as a one-time decision, in fact it is a gradual process of accumulating missed payments on the part of the borrower followed by a lender’s (endogenous) decision to foreclose. Our baseline specification uses three missed payments, but Table 4a shows that we obtain similar estimates of α when defining default as one, two, four, or five missed payments, or as three missed payments for the subset of borrowers who ultimately begin the foreclosure process.²²

Non-recourse State Our results are also not sensitive to whether the default occurs in a non-recourse state or the normalization of the change in income. Table 4a (and online Appendix Figure 8) shows that there is no evidence of more strategic default in states with non-recourse mortgage debt, where defaulting may be more financially advantageous. However, this may not be surprising for two reasons: First, Ghent and Kudlyak (2011) reports that deficiency judgments against borrowers in recourse states are exceedingly rare. Second, Guiso, Sapienza and Zingales (2013) reports no difference between borrowers in recourse and non-recourse states in their subjective estimates of the probability that lenders would pursue them after a default. Finally, while our primary specification examines the percent drop in income at default, online Appendix Figure 9 shows that we find similar results when normalizing the change in income by the monthly mortgage payment due one year before default.

Bank Account Balances We find similar results when looking at bank account balances rather than bank account income. Figure 2a shows that bank account balances deteriorate prior to mortgage default by a similar amount for above and underwater defaulters. There is a modest decline in months 12 to 6 prior to default, followed by a sharp decline in the 6 months immediately prior to default. This time pattern echoes the evolution of income shown in Figure 1a. In addition, the fact that underwater and above water defaulters’ balances deteriorate by a similar amount is further evidence against the presence of strategic default.

Time Period The finding of little strategic default also does not depend on the time

²¹The interpretation of our results is the same if some of the income increases reflect liquidation-inducing expense shocks. An alternative scenario is to consider expense shocks that are unobserved, in the sense that they have no effect on bank account income. The interpretation of our results is unchanged if such unobserved shocks have the same relative prevalence for underwater and above water borrowers as observed income shocks. However, if such shocks are more (resp., less) common for underwater borrowers, then our estimates will understate (resp., overstate) the prevalence of strategic default. Finding appropriate proxies for expenditure shocks is a useful direction for future research.

²²See online Appendix Figures 6 and 7 for the evolution of income around default for each of these subsamples.

period we study, and in particular holds for times when the housing market experienced the most severe distress. Although income data are not available prior to October 2012, bank account balance data are available beginning in January 2007. This means that we have a year of bank account balance history for defaults which occurred in January 2008 or later. Years prior to 2013 are particularly interesting because the mortgage delinquency rate peaked in the first quarter of 2010 and the Case-Shiller house price index bottomed out in early 2012. Figure 2b shows that balance declines are similar for above and underwater borrowers in every year from 2008 through 2014. Likewise, Table 4b shows that the estimates of α are similar across different years.

Loan-to-Value Ratio Finally, our results are also not sensitive to the choice of LTV threshold. Figure 3 shows that the income drop for defaulters is similar across the LTV distribution. Table 4a shows that we obtain similar estimates of α when we define above water borrowers as those with $LTV < 60$ or $LTV < 80$. This allays concerns that measurement error in LTV may lead to misclassification of borrowers who are truly underwater as being above water.

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, we use the bank account balance data, which has a larger sample size and therefore allows us to look even deeper into the LTV distribution. Figure 4 and Table 4b show estimates of α using the balance data for LTV bins in increments of 20 percent. Even for borrowers with LTV as high as 180 percent, a 95 percent confidence interval rules out strategic default by most borrowers. In contrast, we find suggestive evidence for more strategic default behavior for borrowers with LTV above 180, and statistically significant evidence for those with LTV above 220. However, these groups account for a very small share of defaults: only 1.5 percent of defaulters have LTV above 180, and less than 0.2 percent have LTVs above 220.

Figure 4 compares our estimates 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 measure adverse life events. They find that a third of defaults are strategic beginning at LTV of 130 percent and this share is rising in LTV. Thus, in quantitative terms, we find less evidence of strategic default than BDS. However, in qualitative terms, our findings are consistent with BDS, who explicitly conjecture that improvements in the measurement of life events at the individual level may lead to smaller estimates of strategic default.²³

²³BDS write that if after controlling for regional covariates that “the remaining unobserved liquidity shocks [life events] correlate with house price shocks” then the contribution of life events “will be even larger than

4.4 Are *Any* Borrowers Strategic?

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 *any* strategic default. We have already shown that our method can detect strategic default in the small share of the most deeply underwater borrowers. In this section we demonstrate that our method can also detect strategic default in another sub-population: those that miss three straight mortgage payments. However, as with deeply underwater borrowers, 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), who note that a borrower who decides to strategically default will stop making payments once and for all.²⁴ Therefore, if there is evidence of strategic default, it should manifest within this sub-population. At the same time, missing three straight mortgage payments may simply reflect a borrower facing a severe economic shock. However, we do indeed see evidence of strategic default in this sub-group. Online Appendix Figure 10 shows larger income declines for above water defaulters than for underwater defaulters in the subsample that misses three consecutive mortgage payments. In Table 4a, we estimate that 15 percent of underwater defaults are strategic for this subsample.

Another group that we might expect to 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 1), we can not directly speak to the types of investors found to be more strategic in prior work. In particular, Elul, Payne and Tilson (2019) documents that self-declared investors behave similarly to other borrowers. It is *fraudulent* investors who appear more strategic. These are borrowers who claim to be owner-occupants but who in fact have multiple first liens. We do not observe first liens outside of our data set and so cannot identify such borrowers. However, because this sub-population is small, we note that our quantitative estimates of the overall share of strategic default are consistent with some strategic default among these

our estimate already suggests.” One way to interpret our estimates jointly is that this omitted variable of life events *is* correlated with house price shocks. Because the JPMCI data have a measure of this omitted variable that can be observed at the individual level, we learn that strategic default is even less prevalent than it appeared to be in prior work.

²⁴Other papers that study consecutive missed payments include Keys et al. (2012), Bradley, Cutts and Liu (2015) Experian and Wyman (2009) and Tirupattur, Chang and Egan (2010). Keys et al. (2012) measure the share of mortgage defaults that transition straight from 60 days past due to 180 days past due in four months, while remaining otherwise current on all non-HELOC revolving debt. We extend the analysis in Keys et al. (2012) forward through 2015 using the CRISM data and show the results in online Appendix Figure 11. We find that 19.6 percent of underwater borrowers meet their definition of sequential missed payments, while 12.3 percent of abovewater borrowers meet this definition. The excess sequential default rate for underwater borrowers is 7.3 percent. If we interpret this as an alternative estimate of the prevalence of strategic default, it falls within the confidence interval of our central estimate.

borrowers. Elul, Payne and Tilson (2019) reports that fraudulent investors account for only 12 percent of all defaults.

4.5 Double-Trigger vs Cash-flow Defaults

Our central result is that very few defaults are strategic. What more can we say about the remaining non-strategic defaults? Our potential outcomes framework provides a structure for considering this question. In particular, it suggests that the remaining defaults fall into two possible groups: cash-flow defaults (where a life event is a necessary and sufficient condition, captured by the first row in Table 1), and double-trigger defaults (where both a life event and negative equity are necessary for default, captured by the second row in Table 1).

As discussed in the recent review article by Foote and Willen (2018), the standard view in the literature is that any underwater default that is not strategic must be double-trigger. Indeed, there is strong evidence to support double-trigger behavior. For example, seminal work by Gerardi et al. (2018) provided the first direct empirical evidence of double-trigger default. They show that, conditional on a life event, the probability of default is higher for borrowers with negative equity. This implies that both triggers are necessary for some defaults, as in the double-trigger theory.²⁵ More generally, double-trigger default is also consistent with evidence showing that negative equity can causally increase borrower default (Mian and Sufi 2011; Bajari, Chu and Park 2008; Palmer 2015; Gupta and Hansman 2019). 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 ubiquity of above water default and foreclosure documented here and in Low (2018) suggests that negative equity is clearly not a necessary condition for *all* defaults. Because all *above water* defaults must be of the “cash-flow” potential outcome type, it is therefore natural to expect that at least some of the *underwater* defaults are also of this type, rather than exclusively the double-trigger type that has been examined in prior work. We think further investigating the decomposition between cash-flow and double-trigger defaults is an exciting topic for future work.

²⁵We use simulations in online Appendix D to demonstrate how this type of double-trigger behavior is consistent both with the absence of strategic default and with higher default rates for borrowers with negative equity.

5 Re-analysis of Survey Data

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

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

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

We re-analyze the prevalence of underwater strategic default in PSID, adding above water defaulters and all underwater borrowers as comparison groups. Our PSID analysis is in the spirit of the previous sections of this paper, but directly applying the method from Section 2 yields statistically imprecise results.²⁶ Instead, we use the *can pay* definition of strategic default to enable comparability to the prior literature. Furthermore, the PSID also enables us to remedy some shortcomings of the JPMCI data by capturing income across all bank accounts (not just bank accounts with the same institution that services the mortgage) and capturing borrowers’ perceived LTV rather than their estimated LTV.

Figure 5 plots the cumulative distribution function of available resources for above water

²⁶Our method is underpowered in the PSID. We are unable to reject the hypothesis that *no* defaults are strategic and also that *all* defaults are strategic. This is not surprising. In general a noisy variable on the left-hand side of a regression creates larger standard errors, but not bias. Our method, which places a noisy measure of treatment on the left-hand side, sacrifices precision to avoid bias. This is not a problem for a large administrative dataset, but it is a problem for the PSID, where there are only 268 mortgage defaulters.

defaults, underwater defaults, and borrowers who are not in default. Consistent with a role for adverse life events, available resources are much lower for defaulters than non-defaulters. In contrast, when we look *within* defaulters to compare above and underwater borrowers, the distribution of available resources appears to be the same. Thus, the results in the PSID echo the conclusions of Figure 1 (and online Appendix Figure 5) in finding very little evidence of strategic default.

The key reason why we find little strategic default when prior work found substantial strategic default is the adjustment for measurement error, rather than differences in data source or differences in the definition of strategic default. For example, online Appendix Figure 12a shows that 37 percent of underwater defaulters meet the *can pay* definition and one might conclude that these 37 percent of underwater defaults are strategic. Yet online Appendix Figure 12a shows that 39 percent of above water defaulters also meet the *can pay* definition in a sample that has no motive to default strategically. The difference between the share of defaults labeled *can pay* for above water versus underwater is not statistically significant (p-value of 0.82). We conclude that income and consumption obligations are difficult to measure at the household level.²⁷ Not accounting for measurement error may lead researchers to understate the importance of life events, as noted by Gyourko and Tracy (2014).

This null finding appears to be a fundamental feature of the joint distribution of available resources and home equity among defaulters rather than a result of one specific definition of mortgage affordability. GHOW also analyzes an alternative definition of mortgage affordability, which examines how many people would “need to reduce consumption below subsistence levels to remain current on their mortgage”. Using this measure of affordability, online Appendix Figure 12a shows that the share of defaults labeled as strategic is quantitatively similar for above and underwater borrowers (p-value of 0.11). Further, online Appendix Figure 13a shows that the entire distribution of $y - c_{\text{subsistence}} - m$ is similar. In addition, online Appendix Figure 12a shows that the share of defaults with an income increase is similar for above and underwater borrowers in the PSID (p-value of 0.41) and in the JPMCI data (p-value of 0.250). Because above and underwater borrowers have the same distribution of available resources and the same change in income around default, we conclude that there is little evidence of strategic default in the PSID.²⁸

²⁷For example, precise measurement of ability to pay at the time of mortgage default may be difficult if the borrower experienced an income shortfall in one month (even though calendar year income was sufficient to cover the annual mortgage payment). It also might be difficult if the borrower underreported consumption, as people are known to do in recall surveys (Passero, Garner and McCully 2014)

²⁸These figures define underwater as LTV greater than 100, to be consistent with the rest of our paper. Online Appendix Figures 12b and 13b present the same analysis using an LTV cutoff of 90, which is the cutoff used in GHOW.

6 Implications for Modeling Mortgage Default

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

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

A secondary benefit of comparing our empirical results to predictions from a structural model is that the model provides a framework for relaxing one of the common assumptions in empirical work. Specifically, in Section 2 we assumed that each of the treatments was binary (e.g., that each household has an income shock, negative home equity, or both). Put otherwise, the potential outcome function $Y(T^*, G)$ uses binary T^* and binary G . Similarly, the prior empirical work discussed in Section 5 seeks to classify borrowers in a binary fashion as either experiencing a life event or not. In contrast, structural models allow for a cash-flow shock with many possible realizations, such that it is possible to generate an agent’s policy function $Y(T^*, G)$ with discrete T^* .²⁹

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.³⁰ The figure shows that among households with LTV between 100 and 120, the income

²⁹The shock is not continuous in the CC model because the simulation method uses Gaussian quadrature.

³⁰This figure evaluates the joint distribution of income, home equity and default for mortgagors with adjustable-rate mortgages. It shows four scenarios with varying levels of income risk and initial yield rates. CC write that the lessons from this figure are “most visible in Panel D”, which is the scenario with high income risk and high initial yield. We use this scenario throughout our analysis. Online Appendix Figure

of defaulters is substantially lower than the income of non-defaulters. CC explains that this pattern emerges because short-term cash flow considerations drive the default decisions of moderately underwater borrowers:

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

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

This novel prediction from the CC model is also apparent when we replicate our empirical specification within the model. This requires two extensions. First, instead of comparing the level of income of defaulters to that of non-defaulters, our approach relies on calculating the within-borrower change in income prior to default. We show how this statistic varies with LTV both in the baseline CC model and in our data in Figure 6a.³¹ 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 6a shows a substantial income drop both at prepayment for above water borrowers and at default for borrowers with LTV between 100 and 120. The intuition for why an income drop precedes both types of mortgage termination in the model is that they both cause borrowers to lose out on an investment that requires upfront liquidity but has long-term positive expected returns. Thus these borrowers are only likely to terminate if they have suffered a liquidity shock. For expositional simplicity, we use the term “default” to describe both types of mortgage termination. However, the central conclusions from this comparison of model to data are unchanged if we only focus on the income drop gradient for the underwater defaulters and ignore the group of above water prepayers.

¹⁴ shows that we can replicate the summary statistics from CC’s Table II panel D.

³¹Figure 6a reports the income change prior to default normalized by the mortgage payment due. The time interval in CC’s model is annual and it takes a few years for a substantial negative permanent income shock to accumulate because CC use Gaussian quadrature (as is conventional in this literature). We therefore calculate the income change in the model over a four-year time horizon to allow for meaningful negative income shock realizations relative to baseline. We normalize both model and data by mortgage payment due because in the CC model average mortgage payments due by defaulters are equal to 38 percent of income, while in our sample the corresponding ratio is 24 percent.

There are two main conclusions 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 6a 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 2, 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 does a good job of capturing 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 6a shows that the model predicts a drop in income (as a share of mortgage payment due) of 15 percent or less, while in the data the drop is five times as large. A similar divergence holds when considering liquid assets. In the model, we find that these defaulting borrowers have mean liquid assets equal to a bit over one *year* of income. This contrasts with the finding in Table 2 that underwater checking account balances are equal to less than two *weeks* of income, just like the above water borrowers. Intuitively, as LTV increases in the model, more borrowers prefer to default because their homes are a bad financial investment (as in Foster and Van Order 1984), regardless of whether they need what CC calls “immediate budget relief”. In contrast, the decisions of borrowers in the data appear less influenced by the value of the house as a financial asset than is expected in the model. Instead, a substantial income drop and exhaustion of assets precedes default even for deeply underwater borrowers.

Despite this divergence between model and data for deeply underwater borrowers, an extension to the baseline model offers a potential reconciliation. In the baseline model, deeply underwater borrowers default in the absence of immediate cash flow motives because the long-term financial benefits are large while the costs are small. But in practice, the costs of defaulting may also be large. For example, defaulting may impose financial costs through reduced access to credit, or non-financial costs due to borrowers’ attachment to their current home, a fear of social stigma, or a moral aversion to default. The CC model builds in the possibility that these costs are important through an extension that allows for a utility cost

of default, which CC calls *Stigma*. However, the paper explains that “the main difficulty with this extension of our model is determining an appropriate value of *Stigma*.” Indeed, prior work has estimated a wide range of default costs, from as low as a 1.5 percent decrease in the constant equivalent-consumption stream to as high as a 70 percent decrease (Kaplan, Mitman and Violante 2017; Hembre 2018; Schelkle 2018; Laufer 2018).

We propose to use the income drop before default as a new moment to discipline this parameter. Intuitively, the reluctance of borrowers to default on a substantially underwater asset in the absence of immediate budgetary pressure is informative about how costly they perceive this default to be. We therefore estimate the utility cost that minimizes the distance between the model’s predicted income declines for underwater defaulters and the income declines we actually observe in the data. The best fit is that defaulting imposes a utility loss equal to a 25 percent decrease in the constant-equivalent consumption stream. This is a very high cost. We note that this utility cost includes behavioral and moral factors; for example, Guiso, Sapienza and Zingales (2013) report that 82 percent of survey respondents believe that strategic default is morally wrong. The high cost is also consistent with emerging evidence that foreclosures substantially damage family outcomes (Diamond, Guren and Tan 2019).

Under this alternative parameterization, the model is able to closely match our new empirical moments. This is shown visually in Figure 6b. When defaulting is costly, borrowers in the model only exercise their default option when they need short-term budgetary relief. This generates income drops before default consistent with our empirical results. Thus, a high cost of default provides a plausible micro-foundation for the behavior we observe empirically within a benchmark optimizing framework.

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.³² 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

³²See e.g., Corbae and Quintin (2015), Mitman (2016), Kaplan, Mitman and Violante (2017), Guren, Krishnamurthy and McQuade (2018), Campbell, Clara and Cocco (2018), Greenwald, Landvoigt and Van Nieuwerburgh (2018), Diamond and Landvoigt (2019), and Garriga and Hedlund (2020).

might seek to target large income drops before default and low assets at the time of default. In online Appendix Tables 2 and 3, we provide moments on the joint distribution of income, assets, and home equity, which may be useful as a target for such models. Incorporating realistic default behavior triggered by adverse events into macroeconomic models is an exciting topic for future work.

7 Conclusion

This paper asks a simple question: why do borrowers default on mortgages? The literature has focused on two candidate triggers for default: negative equity and life events. However, despite longstanding interest, it has remained difficult to distinguish between these triggers in part because it is difficult to precisely measure life events. We propose a new method to address this measurement error problem using a comparison group of borrowers whose default must have been caused by a 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 dataset linking income and mortgage default. Our central finding is that only three percent of defaults are caused exclusively by negative equity, much less than previously thought. Although our finding contrasts sharply with predictions from standard models, we show that it can be rationalized in models with a high private cost of mortgage default. This reconciliation between model and data may provide a blueprint for general macro finance models where borrower default decisions play a central role.

An additional contribution of this paper is that the method itself may be useful during future housing crises. Indeed, we may be at the precipice of such a crisis due to the Covid-19 recession. If house prices fall and unemployment remains high, a future Fed chair may ask again why borrowers are defaulting. Although we find a clear answer to this question during the last crisis, lenders and policymakers may wonder whether borrower behavior is different this time due to changes in the institutional, policy, or social environment. Such changes could affect the options available to borrowers and the costs and benefits of defaulting. Therefore, this time around, our methodology may be helpful in explaining borrower default in real time.

References

- Albanesi, Stefania, Giacomo De Giorgi, and Jaromir Nosal.** 2017. “Credit Growth and the Financial Crisis: A New Narrative.” National Bureau of Economic Research Working Paper 23740, Cambridge, MA.
- Allcott, Hunt, Benjamin B. Lockwood, and Dmitry Taubinsky.** 2019. “Regressive Sin Taxes, with an Application to the Optimal Soda Tax.” *The Quarterly Journal of Economics*, 134(3): 1557–1626.
- Bajari, Patrick, Chenghuan Sean Chu, and Minjung Park.** 2008. “An Empirical Model of Subprime Mortgage Default From 2000 to 2007.” National Bureau of Economic Research Working Paper 14625.
- Bernanke, Ben S.** 2008. “Mortgage Delinquencies and Foreclosures.”
- Bernstein, Asaf.** 2019. “Negative Home Equity and Household Labor Supply.” Social Science Research Network SSRN Scholarly Paper ID 2700781, Rochester, NY.
- Bhutta, Neil, Jane Dokko, and Hui Shan.** 2017. “Consumer Ruthlessness and Mortgage Default during the 2007 to 2009 Housing Bust.” *The Journal of Finance*, 72(6): 2433–2466.
- Black Knight.** 2020. “Mortgage Monitor.” April 2020.
- Black, Sandra E., Paul J. Devereux, Petter Lundborg, and Kaveh Majlesi.** 2019. “Poor Little Rich Kids? The Role of Nature versus Nurture in Wealth and Other Economic Outcomes and Behaviours.” *The Review of Economic Studies*.
- Boar, Corina, Denis Gorea, and Virgiliu Midrigan.** 2017. “Liquidity Constraints in the U.S. Housing Market.” National Bureau of Economic Research Working Paper 23345.
- Bradley, Michael G., Amy Crews Cutts, and Wei Liu.** 2015. “Strategic Mortgage Default: The Effect of Neighborhood Factors.” *Real Estate Economics*, 43(2): 271–299.
- Campbell, John Y., and João F. Cocco.** 2015. “A Model of Mortgage Default.” *The Journal of Finance*, 70(4): 1495–1554.
- Campbell, John Y., Nuno Clara, and João F. Cocco.** 2018. “Structuring Mortgages for Macroeconomic Stability.” Social Science Research Network SSRN Scholarly Paper ID 3135366, Rochester, NY.
- Campbell, Tim S., and J. Kimball Dietrich.** 1983. “The Determinants of Default on Insured Conventional Residential Mortgage Loans.” *The Journal of Finance*, 38(5): 1569–1581.
- Chan, Sewin, Andrew Haughwout, Andrew Hayashi, and Wilbert Van Der Klaauw.** 2016. “Determinants of Mortgage Default and Consumer Credit Use: The Effects of Foreclosure Laws and Foreclosure Delays.” *Journal of Money, Credit and Banking*, 48(2-3): 393–413.
- Chodorow-Reich, Gabriel.** 2014. “The Employment Effects of Credit Market Disruptions: Firm-level Evidence from the 2008–9 Financial Crisis.” *The Quarterly Journal of Economics*, 129(1): 1–59.
- Christie, Les.** 2010. “How foreclosure impacts your credit score.” *CNN Money*.
- Corbae, Dean, and Erwan Quintin.** 2015. “Leverage and the Foreclosure Crisis.” *Journal of Political Economy*, 123(1): 1–65.
- Dávila, Eduardo.** 2016. “Using elasticities to derive optimal bankruptcy exemptions.” *The Review of Economic Studies*, 87(2): 780–913.
- DeFusco, Anthony, and John Mondragon.** 2018. “No Job, No Money, No Refi: Frictions to Refinancing in a Recession.” Social Science Research Network SSRN Scholarly Paper ID 3231596, Rochester, NY.

- Deng, Yongheng, John M. Quigley, and Robert Van Order.** 2000. "Mortgage Terminations, Heterogeneity and the Exercise of Mortgage Options." *Econometrica*, 68(2): 275–307.
- Diamond, Rebecca, Adam Guren, and Rose Tan.** 2019. "The Effect of Foreclosures on Homeowners, Tenants, and Landlords." *Mimeo*, 35.
- Diamond, William, and Tim Landvoigt.** 2019. "Credit Cycles with Market Based Household Leverage." Social Science Research Network SSRN Scholarly Paper ID 3318481, Rochester, NY.
- Dobkin, Carlos, Amy Finkelstein, Raymond Kluender, and Matthew J. Notowidigdo.** 2018. "The Economic Consequences of Hospital Admissions." *American Economic Review*, 108(2): 308–352.
- Eberly, Janice, and Arvind Krishnamurthy.** 2014. "Efficient Credit Policies in a Housing Debt Crisis." *Brookings Papers on Economic Activity*, 49(2): 73–136.
- Elul, Ronel, Aaron Payne, and Sebastian Tilson.** 2019. "Owner-Occupancy Fraud and Mortgage Performance." Federal Reserve Bank of Philadelphia Working Paper 19-53.
- Epperson, James F., James B. Kau, Donald C. Keenan, and Walter J. Muller.** 1985. "Pricing Default Risk in Mortgages." *Real Estate Economics*, 13(3): 261–272.
- Experian, and Oliver Wyman.** 2009. "Experian–Oliver Wyman market intelligence report: understanding strategic default in mortgages. Part I."
- Fagereng, Andreas, Magne Mogstad, and Marte Ronning.** 2018. "Why Do Wealthy Parents Have Wealthy Children?" Social Science Research Network SSRN Scholarly Paper ID 3146928, Rochester, NY.
- Fannie Mae.** 2011. "Selling Guide."
- Farrell, Diana, Kanav Bhagat, and Chen Zhao.** 2018. "Falling Behind: Bank Data on the Role of Income and Savings in Mortgage Default." JPMorgan Chase Institute.
- Farrell, Diana, Kanav Bhagat, Peter Ganong, and Pascal Noel.** 2017. "Mortgage Modifications after the Great Recession: New Evidence and Implications for Policy." JPMorgan Chase Institute.
- Foote, Christopher L., and Paul S. Willen.** 2018. "Mortgage-Default Research and the Recent Foreclosure Crisis." *Annual Review of Financial Economics*, 10(1): 59–100.
- Foote, Christopher L., Kristopher Gerardi, and Paul S. Willen.** 2008. "Negative equity and foreclosure: Theory and evidence." *Journal of Urban Economics*, 64(2): 234–245.
- Foster, Chester, and Robert Van Order.** 1984. "An Option-Based Model of Mortgage Default." *Housing Finance Review*, 3: 351–372.
- Ganong, Peter, and Pascal Noel.** 2018. "Liquidity vs. Wealth in Household Debt Obligations: Evidence from Housing Policy in the Great Recession." National Bureau of Economic Research Working Paper 24964.
- Garriga, Carlos, and Aaron Hedlund.** 2020. "Mortgage Debt, Consumption, and Illiquid Housing Markets in the Great Recession." *American Economic Review*, 110(6): 1603–1634.
- Gerardi, Kristopher, Kyle F. Herkenhoff, Lee E. Ohanian, and Paul S. Willen.** 2015. "Can't Pay or Won't Pay? Unemployment, Negative Equity, and Strategic Default." National Bureau of Economic Research Working Paper 21630.
- Gerardi, Kristopher, Kyle F. Herkenhoff, Lee E. Ohanian, and Paul S. Willen.** 2018. "Can't Pay or Won't Pay? Unemployment, Negative Equity, and Strategic Default." *The Review of Financial Studies*, 31(3): 1098–1131.

- Ghent, Andra C., and Marianna Kudlyak.** 2011. "Recourse and Residential Mortgage Default: Evidence from US States." *The Review of Financial Studies*, 24(9): 3139–3186.
- Giacoletti, Marco.** 2017. "Idiosyncratic Risk in Housing Markets." Social Science Research Network SSRN Scholarly Paper ID 2995323, Rochester, NY.
- Goldberger, Arthur S.** 1984. "Reverse Regression and Salary Discrimination." *The Journal of Human Resources*, 19(3): 293–318.
- Goldberg, Lawrence, and Charles A. Capone.** 1998. "Multifamily Mortgage Credit Risk: Lessons From Recent History." *Cityscape*, 4(1): 93–113.
- Goodman, Laurie S., Roger Ashworth, Brian Landy, and Ke Yin.** 2010. "Negative Equity Trumps Unemployment in Predicting Defaults." *The Journal of Fixed Income*, 19(4): 67–72.
- Gopalan, Radhakrishnan, Barton H. Hamilton, Ankit Kalda, and David Sovich.** 2019. "Home Equity and Labor Income: The Role of Constrained Mobility." Social Science Research Network SSRN Scholarly Paper ID 2963114, Rochester, NY.
- Greenwald, Daniel, Tim Landvoigt, and Stijn Van Nieuwerburgh.** 2018. "Financial Fragility with SAM?" Social Science Research Network SSRN Scholarly Paper ID 3069621, Rochester, NY.
- Gross, Tal, and Matthew J. Notowidigdo.** 2011. "Health insurance and the consumer bankruptcy decision: Evidence from expansions of Medicaid." *Journal of Public Economics*, 95(7): 767–778.
- Gross, Tal, Raymond Kluender, Feng Liu, Matthew Notowidigdo, and Jialan Wang.** 2019. "Economic Consequences of Bankruptcy Reform." Mimeo.
- Guiso, Luigi, Paola Sapienza, and Luigi Zingales.** 2013. "The Determinants of Attitudes toward Strategic Default on Mortgages." *Journal of Finance*, 68(4): 1473–1515.
- Gupta, Arpit, and Christopher Hansman.** 2019. "Selection, Leverage and Default in the Mortgage Market." *Working Paper*.
- Gupta, Arpit, Edward R. Morrison, Catherine Fedorenko, and Scott Ramsey.** 2017. "Leverage, Default, and Mortality: Evidence from Cancer Diagnoses." Social Science Research Network SSRN Scholarly Paper ID 2583975, Rochester, NY.
- Guren, Adam M.** 2018. "House Price Momentum and Strategic Complementarity." *Journal of Political Economy*, 126(3): 1172–1218.
- Guren, Adam M, Arvind Krishnamurthy, and Timothy J McQuade.** 2018. "Mortgage Design in an Equilibrium Model of the Housing Market." National Bureau of Economic Research Working Paper 24446.
- Gyourko, Joseph, and Joseph Tracy.** 2014. "Reconciling theory and empirics on the role of unemployment in mortgage default." *Journal of Urban Economics*, 80(C): 87–96.
- Hembre, Erik.** 2018. "HAMP, Home Attachment, and Mortgage Default." Working Paper.
- Himmelstein, David U., Elizabeth Warren, Deborah Thorne, and Steffie Woolhandler.** 2005. "Illness and injury as contributors to bankruptcy." *Health Affairs (Project Hope)*, Suppl Web Exclusives: W5–63–W5–73.
- Hsu, Joanne W., David A. Matsa, and Brian T. Melzer.** 2018. "Unemployment Insurance as a Housing Market Stabilizer." *American Economic Review*, 108(1): 49–81.

- Indarte, Sasha.** 2019. “The Impact of Debt Relief Generosity and Liquid Wealth on Household Bankruptcy.” Social Science Research Network SSRN Scholarly Paper ID 3378669, Rochester, NY.
- Kaplan, Greg, Kurt Mitman, and Giovanni L Violante.** 2017. “The Housing Boom and Bust: Model Meets Evidence.” National Bureau of Economic Research Working Paper 23694.
- Keys, Benjamin J., Tomasz Piskorski, Amit Seru, and Vikrant Vig.** 2012. “Mortgage Financing in the Housing Boom and Bust.” In *Housing and the Financial Crisis*. 143–204.
- Kleven, Henrik, Camille Landais, and Jakob Egholt Sogaard.** 2020. “Does Biology Drive Child Penalties? Evidence from Biological and Adoptive Families.” National Bureau of Economic Research Working Paper 27130.
- Laufer, Steven.** 2018. “Equity Extraction and Mortgage Default.” *Review of Economic Dynamics*, 28: 1–33.
- Low, David.** 2018. “Mortgage Default with Positive Equity.” *Working Paper*.
- Mayer, Christopher, Edward Morrison, Tomasz Piskorski, and Arpit Gupta.** 2014. “Mortgage Modification and Strategic Behavior: Evidence from a Legal Settlement with Countrywide.” *American Economic Review*, 104(9): 2830–2857.
- Mian, Atif, and Amir Sufi.** 2011. “House Prices, Home Equity-Based Borrowing, and the US Household Leverage Crisis.” *American Economic Review*, 101(5): 2132–2156.
- Mitman, Kurt.** 2016. “Macroeconomic Effects of Bankruptcy and Foreclosure Policies.” *American Economic Review*, 106(8): 2219–2255.
- Office of the Comptroller of the Currency.** 2005. “CREDIT RISK MANAGEMENT GUIDANCE FOR HOME EQUITY LENDING.”
- O’Malley, Terry.** 2018. “The Impact of Repossession Risk on Mortgage Default.” Central Bank of Ireland Research Technical Papers 1/RT/18.
- Palmer, Christopher.** 2015. “Why Did So Many Subprime Borrowers Default During the Crisis: Loose Credit or Plummeting Prices?” *Working Paper*.
- Passero, William, Thesia I. Garner, and Clinton McCully.** 2014. “Understanding the Relationship: CE Survey and PCE.” In *Improving the Measurement of Consumer Expenditures*. 181–203. University of Chicago Press.
- Pearl, Judea.** 1999. “Probabilities Of Causation: Three Counterfactual Interpretations And Their Identification.” *Synthese*, 121(1): 93–149.
- Ridder, Geert, and Robert Moffitt.** 2007. “The Econometrics of Data Combination.” Elsevier Handbook of Econometrics.
- Riddiough, Timothy J.** 1991. “Equilibrium mortgage default pricing with non-optimal borrower behavior.” PhD diss. Univseristy of Wisconsin.
- Riddiough, Timothy J., and Steve B. Wyatt.** 1994. “Strategic default, workout, and commercial mortgage valuation.” *The Journal of Real Estate Finance and Economics*, 9(1): 5–22.
- Rosenbaum, Paul R.** 2001. “Effects Attributable to Treatment: Inference in Experiments and Observational Studies with a Discrete Pivot.” *Biometrika*, 88(1): 219–231.
- Roubini, Nouriel.** 2008. “The Forthcoming “Jingle Mail” Tsunami: 10 to 15 Million Households Likely to Walk Away from their Homes/Mortgages Leading to a Systemic Banking Crisis.”

- Rubin, Donald B.** 1974. “Estimating causal effects of treatments in randomized and nonrandomized studies.” *Journal of Educational Psychology*, 66(5): 688–701.
- Scharlemann, Therese, and Stephen H. Shore.** 2019. “The Effect of Mortgage Payment Size on Default and Prepayment: Evidence from HAMP Resets.” *Working Paper*.
- Schelkle, Thomas.** 2018. “Mortgage Default during the U.S. Mortgage Crisis.” *Journal of Money, Credit and Banking*, 50(6): 1101–1137.
- Tirupattur, V, O Chang, and J Egan.** 2010. “Understanding strategic defaults.” Morgan Stanley Research.
- Vandell, Kerry D.** 1995. “How Ruthless Is Mortgage Default? A Review and Synthesis of the Evidence.” *Journal of Housing Research*, 6(2): 245–264.
- Wooldridge, Jeffrey M.** 2010. *Econometric Analysis of Cross Section and Panel Data*. Vol. 1 of *MIT Press Books*, The MIT Press.
- Yamamoto, Teppei.** 2012. “Understanding the Past: Statistical Analysis of Causal Attribution.” *American Journal of Political Science*, 56(1): 237–256.

Figure 1: Income in Year Prior to Mortgage Default

Notes: This figure describes income in the year prior to mortgage default in the JPMCI data. In panel (a), the squares show mean income of underwater borrowers with negative home equity in comparison to two benchmarks. These benchmarks capture two different motivations for mortgage default. The first benchmark, which captures the negative equity (strategic) motive, shows average income for all underwater borrowers in circles. The second benchmark in triangles shows the path of income for above water borrowers with positive home equity; because this group has no strategic default motive, their defaults are triggered by a life event that reduced their cash-flow. Mean income is normalized by average initial income, which is computed separately for above and underwater borrowers. Panel (b) shows the full distribution of income change from the month of default relative to 12 months prior by home equity and is truncated at -1 and 1 to improve readability. Default is defined as three missed mortgage payments. See Section 4 for details.

Figure 2: Bank Account Balance in Year Prior to Mortgage Default

Notes: This figure shows the change in bank account balances in the year prior to mortgage default by home equity in the JPMCI data. The JPMCI data include monthly bank account balances beginning in January 2007 and monthly income beginning in October 2012. The top panel shows the monthly evolution of balances for the same income sample shown in the prior figures. The bottom panel shows the evolution of balances separately for each year, using the full balance sample that goes back to 2007 (allowing us to measure one year of balance change for defaults starting in 2008). Default is defined as three missed mortgage payments.

Figure 3: Income Drop Before Mortgage Default by Loan-to-value Ratio

Notes: This figure describes the change in income in the twelve months prior to mortgage default in the JPMCI data. Vertical lines indicate the 95 percent confidence intervals.

Figure 4: Defaults Causally Attributable to Life Events: Heterogeneity by Loan-to-value Ratio

Notes: This figure compares estimates of the share of defaults that are causally attributable to life events in Bhutta, Dokko and Shan (2017) and in the JPMCI bank account data. The bank account estimates use equation (3) with data on the change in bank account balances prior to default. The Bhutta, Dokko and Shan (2017) estimates come from Figure 6 of that paper, where the proportion of defaults that are causally attributable to life events is 100 percent minus the share of defaults that are “strategic”. The vertical lines show 95 percent confidence intervals for our estimates. Online Appendix Figure 15 replicates the same analysis using the change in income prior to default. The share of defaulters with high LTVs is small: 1.0 percent of defaulters have LTV of 180-200, 0.3 percent have 200-220, and 0.2 percent have LTV 220+.

Figure 5: Distribution of Available Resources is Same for Above Water and Underwater Defaulters in PSID

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

Figure 6: Income Drop at Default Compared to Prior Theoretical Literature

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

Table 1: Potential Outcomes Model of the Mortgage Default Decision

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

Notes: This table shows the different combination of potential outcomes from the environment described in Section 2. We assume that borrowers without a life event or negative equity never default ($Y(0,0) = 0$). There are therefore $2^3 = 8$ possible combinations for the remaining three binary potential outcomes. However three rows with a 1 in the columns for $Y(1,0)$ or $Y(0,1)$ and a 0 in the $Y(1,1)$ column are ruled out by monotonicity. Furthermore, we focus on defaulters, and so also leave out the “never default” type where all potential outcomes are zero regardless of life events or negative equity. This leaves the four rows shown. The α parameter identified in equation (3) quantifies the share of underwater defaults accounted for by the top two rows, highlighted in gray.

Table 2: Summary Statistics by Home Equity

Variable	Above water	Underwater
Combined loan-to-value ratio (%)	71	121
Monthly bank account income (\$)	4,053	4,436
Bank account balance (\$)	1,455	1,692
Property value (\$)	243,094	235,149
Monthly mortgage payment due (\$)	966	1,137
Age	50	50
Share with joint deposit account	0.4	0.44
N	22,687.0	6,347.00

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

Table 3: Income Drop at Default by Home Equity

	<i>Dependent variable:</i>			
	Change in income from one year before default			
	Mean	Median	p25	p75
	(1)	(2)	(3)	(4)
Date of default	−0.203 (0.004)	−0.175 (0.005)	−0.145 (0.004)	−0.269 (0.008)
Date of default * underwater	0.006 (0.007)	−0.014 (0.010)	−0.004 (0.010)	−0.006 (0.017)
N mortgages	29,034	29,034	29,034	29,034
Observations	174,204	174,204	174,204	174,204

Notes: This table describes the income drop at default in the JPMCI data. The sample includes a pre-period (12 to 10 months prior to default) and a default period (2 months prior to default through the month of default). The dependent variable is the ratio of monthly income to average income in the pre-period. The regression specification is $\frac{Income_t}{Income_{pre}} = \lambda + \kappa(LTV > 100) + \gamma \mathbf{1}(t = -2, -1, 0) + \beta \mathbf{1}(t = -2, -1, 0) \times LTV > 100 + \varepsilon$, as shown in equation (5). The table reports estimates for the change in income during the default period ($\hat{\gamma}$) and the interaction with being underwater ($\hat{\beta}$). Column (1) analyzes the mean change in income. Columns (2), (3) and (4) show the 50th, 25th and 75th percentiles of the change in income respectively. 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.

Table 4: Share of Defaults Causally Attributable to Life Events ($\hat{\alpha}$)

(a) Estimates Using Income Data

Category	$\hat{\alpha}$ (SE)	P-value vs baseline
Baseline	0.97 (0.04)	-
Percentiles		
25th	0.97 (0.08)	0.995
50th	1.07 (0.06)	0.128
75th	1.01 (0.03)	0.393
Abovewater as LTV <60	0.99 (0.04)	0.679
Abovewater as LTV <80	1.02 (0.05)	0.458
LTV		
101-120	0.97 (0.05)	0.997
121-140	0.95 (0.07)	0.791
141-160	0.90 (0.11)	0.567
161+	1.15 (0.15)	0.245
Three consecutive missed payments	0.85 (0.04)	0.016
Non recourse states	0.95 (0.07)	0.745
Foreclosure start	0.85 (0.06)	0.074
Months past due		
1	1.06 (0.10)	0.372
2	0.90 (0.05)	0.287
4	0.98 (0.06)	0.911
5	1.01 (0.07)	0.598

(b) Estimates Using Balance Data

Category	$\hat{\alpha}$ (SE)	P-value vs baseline
Baseline	0.99 (0.03)	-
Year		
2008	1.13 (0.10)	0.163
2009	0.92 (0.05)	0.235
2010	0.97 (0.07)	0.844
2011	0.91 (0.04)	0.117
2012	1.09 (0.07)	0.178
2013	1.10 (0.10)	0.289
2014	0.93 (0.12)	0.630
2015	0.98 (0.28)	0.970
LTV		
101-120	1.01 (0.03)	0.609
121-140	0.97 (0.04)	0.795
141-160	0.96 (0.06)	0.647
161-180	1.05 (0.09)	0.482
181-200	0.86 (0.11)	0.285
201-220	0.77 (0.20)	0.295
221+	0.30 (0.28)	0.014

Notes: This table reports estimates of $\hat{\alpha}$, which is the share of defaults causally attributable to life events, in the JPMCI data. $\hat{\alpha}$ is constructed using equation (3). See Sections 4.2 and 4.3 for details.

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

Peter Ganong and Pascal Noel

Contents

A	Figures and tables	1
B	Data appendix	16
C	Proof of proposition 1	16
D	Double-trigger simulations	18

A Figures and tables

Figure 1: Prevalence of Above Water Mortgage Default

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

Figure 2: Payment Due And Payment Made Prior to Default

Notes:

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

Figure 3: Social Security Income Change by Home Equity

Notes: This figure shows the change in Social Security income in the year prior to default in the JPMCI data.

Figure 4: Seasonality in Mortgage Default

Notes: This figure documents that correlated seasonality in mortgage default and in income generates an S-shape in the “all underwater borrowers” series in Figure 1a. Panel (a) shows the distribution of default dates. In unreported results, we find similar seasonal patterns in broader samples in the JPMCI data and also in the CRISM data. Panel (b) shows mean post-tax income by calendar month. This pattern, too, also holds in broader samples in the JPMCI data. Panel (c) shows the role that correlated seasonality plays in generating the S-shape in the “all underwater borrowers” series. The purple circles replicate the series from Figure 1a. That figure uses equation (4), which indexes time as $t \in \{Oct2013, Nov2013 \dots Aug2015\}$. The blue squares re-estimate equation (4), where $t \in \{Jan, Feb \dots Dec\}$. This procedure isolates the effect of seasonal patterns on the “all underwater borrowers” series by capturing seasonality in defaults (shown in panel a) and seasonality in income (shown in panel b). Because defaults are highest at the end of each calendar year (panel a) and income is highest in March and April around tax refund season and December around Christmas (panel b), the “all underwater borrowers” series has peaks in the month of default and roughly eight months prior.

Figure 5: Distribution of Income Change Prior to Mortgage Default

Notes: This figure shows the cumulative distribution function for the change in income, divided by average initial income. Average initial income is computed one year prior to mortgage default and is computed separately for underwater and above water borrowers. The ratio can be less than -1 or more than 1 since some households have income declines or increases that is larger than the average income level. This figure provides an alternative visualization of the histogram in Figure1b.

Figure 6: Income by Alternative Missed Payment Thresholds and Home Equity

(a) One Month Past Due

(b) Two Months Past Due

(c) Four Months Past Due

(d) Five Months Past Due

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

Figure 7: Evolution of Income by Home Equity – Foreclosure Sample

Notes: This figure replicates Figure 1a for the subset of defaulters who ultimately begin the foreclosure process.

Figure 8: Income Prior to Default by State-level Availability of Recourse

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

Figure 9: Income Change as Share of Payment Due by Home Equity

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

Figure 10: Income Prior to Default by Consecutive Missed Payments

Notes: This figure replicates Figure 1a from the JPMCI data for the subset of borrowers who miss three consecutive payments. Borrowers who miss three consecutive payments are 48 percent of underwater defaults and 39 percent of above water defaults.

Figure 11: Share of Mortgage Defaults with Consecutive Missed Payments

Notes: This figure extends the analysis in Keys et al. (2012) using the CRISM data. That paper measures the share of mortgage defaults that transition straight from 60 days past due to 180 days past due in four months, while remaining otherwise current on all non-HELOC revolving debt. We refer to such defaults as “straight and otherwise current”. The average share of defaults that meet this definition is 19.6 percent of defaults for underwater borrowers and 12.3 percent of defaults for above water borrowers. Thus, the excess share of straight and otherwise current defaults for underwater borrowers is 7.3 percent.

Figure 12: Alternative Measures of Strategic Default

(a) Share of Defaults

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

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

Figure 13: Alternative Measures of Strategic Default – Distributions

(a) Available Resources Using Subsistence Measure

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

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

Figure 14: Campbell and Cocco (2015) Structural Model Replication

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

Figure 15: Defaults Causally Attributable to Life Events: Heterogeneity by Loan-to-value Ratio

Notes: This figure compares estimates of the share of defaults that are causally attributable to life events in Bhutta, Dokko and Shan (2017) and in the JPMCI bank account data. The bank account estimates use equation (3) with data on the change in income prior to default. The Bhutta, Dokko and Shan (2017) estimates come from Figure 6 of that paper, where the proportion of defaults that are causally attributable to life events is 100 percent minus the share of defaults that are “strategic”. This figure replicates Figure 4 using income data (whereas Figure 4 uses balance data).

Figure 16: Double Trigger Is Consistent with Zero Strategic Default

(a) Real Data: Regress Y on T

(b) Simulated Data: Regress Y on T

(c) Simulated Data: $E(T|Y = 1)$

Notes: The top panel computes binned default rates for 8 quantiles of the normalized change in income by home equity using the PSID. The probability of default is higher for underwater borrowers with a decline in income than it is for above water borrowers with a decline in income. We use a simulation to show that this finding is consistent with a finding of no strategic default. The middle and bottom panel show two different ways of visualizing the relationship between income, home equity, and default. Panel (b) replicates the analysis in panel (a), while panel (c) replicates the analysis in Figure 1a. See online Appendix D for details.

Table 1: Summary Statistics: Chase versus CRISM and McDash

Sample	Benchmark	Chase	CRISM	McDash
All mortgages	90 day delinquency rate	3.2%	3.3%	3.8%
All mortgages	Share investor	6.8%	3.9%	5.6%
All mortgages	Share primary occupant	89%	93%	91%
All mortgages	Share underwater	19%	22%	
Defaulters	Share investor	6.4%	4.3%	5.9%
Defaulters	Share primary occupant	90%	94%	92%
Defaulters	Share underwater	50%	58%	
Defaulters	Share of above water defaults with foreclosure within year	40%	55%	
Defaulters	Share of underwater defaults with foreclosure within year	45%	57%	

Notes: This table compares summary statistics regarding the matched mortgage-bank account dataset from Chase to the Credit Risk Insight Servicing McDash (CRISM) dataset in 2011. The CRISM dataset is constructed by linking credit bureau records from Equifax with mortgage servicing records from McDash. We use 2011 as the comparison year because this is the year when US house prices reached their nadir. Investor and primary occupant are reported by borrowers at mortgage origination. “Foreclosure” indicates that the mortgage servicer initiated foreclosure proceedings.

Table 2: Income and Assets of Defaulters by Loan-to-Value

LTV	Drop as share of income	Drop as share of mortgage payment due	Checking Balance
<100	-0.203	-0.878	1,038
100-120	-0.198	-0.795	1,159
120-140	-0.195	-0.785	1,234
140-160	-0.185	-0.707	1,104
160+	-0.235	-0.953	1,080

Notes: This table measures economic conditions at the time of default by loan-to-value (LTV) bin. The first two columns show measures of the average income drop from one year prior to default to the month of default and the third column shows mean checking account balances at the date of default. Note that this table describes borrowers at the date of default, which is different from Table 2 in the main text, which describes borrowers one year before default.

Table 3: Distribution of Checking Account Balances of Defaulters

LTV	p10	p25	p50	p75	p90
Above water	-0.6	22.8	277.8	1,118.8	2,610.6
Underwater	0.4	31.4	358.1	1,301.4	2,939.5

Notes: This table shows the distribution of checking account balances at the date of default for the primary analysis sample in the JPMCI data. To avoid disclosing information for any single household, the table reports pseudo-medians based on cells of at least 10 observations. A negative balance indicates that an account is in overdraft status. Thus, this table shows that about 10 percent of above water and underwater defaulters have overdrafted their checking accounts. Note that this table describes borrowers at the date of default, which is different from Table 2 in the main text, which describes borrowers one year before default.

B Data appendix

To be included in the analysis sample, we require that the household have an open checking account from one year before default through the date of default.

The income data are available from October 2012 forward. Thus, for inclusion in the income analysis sample, we define the date of default as the first date *after October 2012* when a mortgage was 90 days past due.

The unit of observation is a mortgage. There are 29,034 mortgages which meet this definition of default, have reliable data on payments made, have non-missing loan-to-value ratios, and have income data available for one year prior to default. There are 28,589 unique households associated with these 29,034 mortgages; this situation arises because there are a very small number of households that default on multiple first lien mortgages that are serviced by Chase.

We apply the same logic to the checking account balance sample. Balance is measured at the beginning of the month. Inclusion in this sample requires having a checking account open in the year prior to the first default, where the date of default is the first default after January 2007.

C Proof of proposition 1

$$\begin{aligned}
\alpha_{G=1} &\equiv \frac{E(Y|G=1) - E(Y_{01}|G=1)}{E(Y|G=1)} \\
&= 1 - \frac{E(Y_{01}|G=1, T^*=0)}{E(Y|G=1)} \\
&= 1 - \frac{P(Y=1|T^*=0, G=1)}{P(Y=1|G=1)} \\
&= 1 - \frac{P(T^*=0|Y=1, G=1)}{P(T^*=0|G=1)} \tag{7}
\end{aligned}$$

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

The law of iterated expectations implies that

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

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

$$\begin{aligned} P(T^* = 0|Y = 1, G = 1) &= \frac{E(T(1)|T^* = 1, Y = 1, G = 1) - E(T|Y = 1, G = 1)}{E(T(1)|T^* = 1, Y = 1, G = 1) - E(T(0)|T^* = 0, Y = 1, G = 1)} \\ &= \frac{E(T(1)) - E(T|Y = 1, G = 1)}{E(T(1)) - E(T(0))} \end{aligned} \quad (8)$$

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

$$\begin{aligned} E(T(1)) &= E(T|Y = 1, G = 0, T^* = 1)P(T^* = 1|Y = 1, G = 0) \\ &= E(T|Y = 1, G = 0) \end{aligned} \quad (9)$$

where $P(T^* = 1|Y = 1, G = 0) = 1$ by assumption 1. Substitute equation (9) into the numerator of equation (8) to get

$$P(T^* = 0|Y = 1, G = 1) = \frac{E(T|Y = 1, G = 0) - E(T|Y = 1, G = 1)}{E(T(1)) - E(T(0))} \quad (10)$$

This expression captures the numerator of the ratio in equation (7). Applying the same logic to the denominator in the ratio of equation (7) gives

$$\begin{aligned} P(T^* = 0|G = 1) &= \frac{E(T(1)|T^* = 1, G = 1) - E(T|G = 1)}{E(T(1)|T^* = 1, G = 1) - E(T(0)|T^* = 0, G = 1)} \\ &= \frac{E(T(1)) - E(T|G = 1)}{E(T(1)) - E(T(0))} \\ &= \frac{E(T|Y = 1, G = 0) - E(T|G = 1)}{E(T(1)) - E(T(0))} \end{aligned} \quad (11)$$

where $E(T|G = 1)$ includes both underwater defaulters and non-defaulters. We take the ratio of equations (10) and (11). The denominators ($E(T(1)) - E(T(0))$) cancel, so

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

Plugging this ratio into equation (7) gives

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

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

D Double-trigger simulations

Our empirical finding of no strategic default is consistent with this prior evidence in favor of the “double-trigger” theory of default. We use a simple simulation to illustrate this point. Because the model is highly stylized, our results should be taken more as suggestive illustrations of the economic forces at play rather than as an empirically-realistic model parameterization.

Define $Y(T^*, \eta)$ where Y is a binary variable for default, T^* is a binary random variable that is set to one when the household receives a life event shock and η is a second binary random variable (a preference shock) which determines whether the borrower defaults conditional on receiving a life event shock. We model the default decision as

$$Y(T^*, \eta) \equiv T^* \times \eta$$

so a borrower defaults only if they experience a life event and decide to default. Note that there is no “strategic default” in this model by borrowers who “can pay”, to use the language of Gerardi et al. (2018).

We allow the preference shock to vary with household leverage, so the probability of default conditional on a life event can differ between above water and underwater borrowers. We assume that adverse life events T^* are measured with noise. Specifically, we assume the change in observed income T is the sum of T^* and mean-zero noise ε , where ε has a normal distribution with standard deviation σ ($T = T^* + \varepsilon$, $\varepsilon \sim N(0, \sigma)$). We parameterize the model’s three parameters as follows: $P(T^* = 1) = 0.25$, $\sigma^2 = 1.5$, and

$$P(\eta = 1) = \begin{cases} 0.1 & \text{if abovewater} \\ 0.5 & \text{if underwater} \end{cases}.$$

Using actual data, online Appendix Figure 16a shows that among borrowers with a positive income change (who should be able to pay if income is measured without error), the default rate is higher for borrowers who are underwater than for borrowers who are above water.²

Using simulated data, online Appendix Figure 16b plots the mean default rate against the eight bins of the change in income using the simulated data. It shows that the intercept is higher and the slope is steeper for underwater borrowers than for above water borrowers.

Finally, we apply the specification used in Section 2 to the simulated data to show that it is consistent with no strategic default and with our findings in the bank data. Online Appendix Figure 16c shows the change in average income among defaulters, separately for underwater and above water. The average change in income is the same for underwater defaults and above water defaults, just as we find in the bank data. Thus, we show that our finding that a life event is a necessary condition for default is consistent with prior double-trigger finds in the literature.

²This approximately replicates the Gerardi et al. (2018) Table 4 finding that among “can pay” borrowers, the default rate is higher for underwater borrowers than above water. Gerardi et al. (2018) Table 5 use *residual income*, which is “the difference between household resources and the mortgage payment exactly”. We focus here on *income changes* for consistency with our specification in the JPMCI data.