The Journal of FINANCE
The Journal of THE AMERICAN FINANCE ASSOCIATION

THE JOURNAL OF FINANCE • VOL. LXXVI, NO. 2 • APRIL 2021

The Impact of Repossession Risk on Mortgage Default

TERRY O'MALLEY*

ABSTRACT

I study the effect of removing repossession risk on a mortgagor's decision to default. Reducing default costs may result in strategic default, particularly during crises when homeowners can be substantially underwater. I analyze difference-in-differences variation in repossession risk generated by an unexpected legal ruling in Ireland that prohibited collateral enforcement on delinquent residential mortgages originated before a particular date. I estimate that borrowers defaulted by 0.3 percentage points more each quarter after the ruling, a relative increase of approximately one-half. High loan-to-value ratios and low liquidity are associated with a larger treatment effect, suggesting both equity and consumption-based motivations.

DURING ECONOMIC CRISES, GOVERNMENTS often consider debt relief policies that reduce repossession risk for mortgage borrowers, such as foreclosure moratoria, because of the negative externalities and social disruption of widespread evictions.[1] However, these policies come with a potentially large moral hazard cost. By lowering the costs of default, reducing repossession risk makes default more attractive. This association may be particularly strong during recessions, when underwater homeowners have strong incentives to default. Absent repossession risk, some borrowers who would have otherwise continued paying choose to default, substantially reducing the net benefits of

*Terry O'Malley is at the Central Bank of Ireland. I thank Morgan Kelly and Fergal McCann for detailed comments and guidance throughout this project. I also thank Amit Seru and the editorial team, whose comments greatly improved the paper. I acknowledge helpful comments from Benjamin Arold; David Byrne; Paul Devereux; Andreas Fuster; Edward Gaffney; Brian Higgins; Robert Kelly; Marianna Kudlyak; Philip Lane; Michael McMahon; Rachael Meager; Conor O'Toole; Kevin Schnepel; Johannes Stroebel; Karl Whelan; as well as from participants at the 2017 Irish Economic Association Annual Conference, 2017 European Economic Association Annual Congress, 2017 Empirics and Methods in Economics Conference, 2018 European Doctoral Group in Economics Jamboree, and both Central Bank of Ireland and University College Dublin seminars. Special thanks to John Joyce for assistance with the data. I have read *The Journal of Finance* disclosure policy and have no conflicts of interest to disclose.

Correspondence: Terry O'Malley, Central Bank of Ireland, Dublin 1, Ireland; e-mail: terry.omalley@ucd.ie.

[1] A recent literature from the United States foreclosure crisis shows large negative externalities of foreclosures in terms of house prices (Campbell, Giglio, and Pathak (2011), Mian, Sufi, and Trebbi (2015)), underinvestment in the housing stock (Melzer (2017)), and aggregate demand externalities (Mian, Sufi, and Trebbi (2015)).

DOI: 10.1111/jofi.12990

© 2020 the American Finance Association

the policy. In this paper, I estimate the magnitude of the impact of removing repossession risk on default by examining a recent natural experiment that occurred in Ireland.

Clean estimates of the moral hazard costs of debt relief policies such as foreclosure moratoria are difficult to obtain. Since restricting debt relief to a particular group of borrowers partially negates its benefits, natural experiments are rare. Moreover, while foreclosure moratoria were widely used across the United States during the Panic of 1819 and the Great Depression (Alston (1984), Skilton (1943)) and internationally during the recent Great Recession albeit less extensively (Artavanis and Spyridopoulos (2018), Gabriel, Iacoviello, and Lutz (2020)), to date there have been few situations in which both the moratoriums apply differentially to similar borrowers and the necessary data are available.

Ireland, however, offers an interesting setting to study the moral hazard cost of debt relief policies. Owing to a poorly drafted change in Irish law governing property transfers and its later unexpected discovery during a court case, the repossession regime of existing mortgage contracts was retroactively amended. Specifically, in July 2011, a judge ruled that properties mortgaged before December 2009 could no longer be repossessed in the event of default.

To estimate the impact of repossession risk on mortgage default, I compare the performance of two loan vintages separated by the December 2009 cutoff date, both before and after the July 2011 legal ruling (known as the "Dunne judgment"). Using regulatory panel data, I construct a sample of loans originated over a 180-day window around the cutoff date, and then, match treatment and control groups to obtain the difference-in-differences (DD) variation in repossession risk. The two groups are similar on observables before the Dunne judgment, consistent with the exogeneity of the origination cutoff date. However, while the groups follow parallel trends before the ruling, their default rates diverge markedly afterwards, with the largest effect observed in the quarter immediately following the ruling.

Specifically, I estimate that the Dunne judgment increased the quarterly default rate by 0.3 percentage points. In the control group, a 0.3 percentage point higher default rate is associated with variation between the median loan-to-value (LTV) ratio plus 38 percentage points, or the median interest rate plus two percentage points. When scaled by the estimated counterfactual, this translates to a relative effect of between 40% and 60%.[2]

To gauge the robustness of this result, I estimate a variety of models under different assumptions about confounding factors. Alongside a straightforward DD regression, I consider specifications with additional borrower and loan covariates and with high-dimensional fixed effects. Since treatment status is determined by a loan's origination date, loan vintage is a nuisance parameter under my research design. I therefore also estimate models that adjust for potentially confounding cohort effects by including vintage-calendar time parameters in the estimation. The significance of the treatment effect is robust to these different specifications. Graphical evidence is also consistent with

[2] Table II contains a range of absolute effects.

excess default only for loans issued before the cutoff date, although a regression discontinuity specification fails to rule out a null treatment effect. In addition, the treatment effect remains statistically significant in a permutation exercise in which I simulate the null distribution of the treatment effect with 1,000 placebo analyses.

Standard models of mortgage default, known as "double-trigger" models, demonstrate that borrowers will default if the intrinsic financial value of their mortgage does not exceed the costs of defaulting.[3] In line with this view, recent empirical research shows that large costs of default are required to match the aggregate or available micro data from the U.S. default crisis (Gerardi et al. (2017), Laufer (2018), Schelkle (2018)). Substantial default costs therefore act as a constraint on mortgage defaults during a recession, when prices and/or incomes fall.[4] Default costs are both pecuniary and nonpecuniary in nature (Guiso, Sapienza, and Zingales (2013)). Pecuniary costs include the current rent that a borrower would have to pay if she were evicted from her home as well as the loss of a permanent hedge against future fluctuations in rent. Nonpecuniary costs include the loss of utility derived from a customized home, local amenities, and social networks. A borrower experiences all of these costs from loss of ownership. Removing or substantially reducing the risk of repossession may therefore increase delinquency rates during a crisis, when the financial value of mortgages may be substantially diminished.

The main contribution of this paper is to provide clean estimates of the default response of borrowers to a large reduction in one of the most important costs of default, namely, repossession risk. Empirical estimates of the sensitivity of strategic behavior to default costs are important for understanding the costs of not only foreclosure moratoria, but also of debt relief policies that influence default costs more generally. Evidence on the effects of such programs is particularly important in light of the experience of the Great Recession, when efforts to alleviate household financial crises were impeded by financial-intermediation frictions.[5]

[3] The intrinsic value of a mortgage comprises the option value of repayment, which reflects the expected investment profit from potential positive equity in the future and the ability to default in the future, minus the mortgage payment amount. Campbell and Cocco (2015) develop a quantitative model of default among borrowers who demonstrate optimizing behavior. Schelkle (2018) and Laufer (2018) provide structural models. See Fuster and Willen (2017) for a straightforward discussion and Foote and Willen (2018) for a review of the historical development of mortgage default models.

[4] Recounting U.S. government deliberations on debt policies in response to the Great Recession, Goolsbee (2014) highlights the importance of default costs in acting as a wedge between the intrinsic value of a mortgage and its value to the homeowner.

[5] A large recent literature has documented various frictions in financial intermediation that dampen the pass-through of monetary easing or hinder endogenous renegotiation. These frictions include a lack of market for negative equity refinancing; barriers to refinancing for the unemployed; frictions in loan modifications arising from information asymmetries, securitization, and organizational constraints (Piskorski, Seru, and Vig (2010), Agarwal et al. (2011), Adelino, Gerardi, and Willen (2013), Hurst et al. (2016), Agarwal et al. (2017), Beraja et al. (2019), DeFusco and Mondragon (2020)).

The closest paper to mine is Mayer et al. (2014), who find a substantial strategic response to a modification program. Agarwal et al. (2017) find no evidence of strategic default in response to the Home Affordable Modification Program (HAMP), but that policy was designed to minimize incentives for strategic behavior. Ghent and Kudlyak (2011) use cross-state variation in recourse status, which is largely due to historical precedent, and find positive impacts of nonrecourse mortgages on default. Collins and Urban (2015), in contrast, find no effect of an eight-month foreclosure moratorium on default.

My data also allow me to estimate the relevant margins of excess default. Identifying the reasons for strategic default is important because policymakers might be willing to accept the moral hazard cost to target debt relief to a specific group. The double trigger model suggests that borrowers default more when default costs fall for two reasons. First, underlower default costs, borrowers default at a lower home-equity threshold. In a purely financial sense, paying the mortgage on an underwater home is rational if the possibility of positive equity in the future justifies the interest payment today. Some borrowers who were previously paying an underwater mortgage did so not because of the option value of repayment, but because of the high costs of default—absent such default costs, the bet on future house prices does not warrant payment. Second, under lower default costs, liquidity-constrained borrowers also default more. When default costs fall, these strategic defaulters choose to no longer sacrifice consumption to pay their mortgage and therefore default on their payments.[6]

To investigate the empirical relevance of these two margins, I analyze heterogeneous treatment effects. For a subset of loans, I link the borrower's liquid account balance at the same bank and estimate triple-difference regressions. In separate models using pretreatment covariates, I find large treatment effects for both highly levered and low-liquidity borrowers. The treatment effect is approximately 1.5 percentage points in the lowest tercile of the liquid wealth distribution and 2 percentage points in the highest tercile of the LTV distribution. Borrowers with high liquid wealth or low LTVs show no pattern of excess default.

Given the above evidence, a second contribution of my paper is to demonstrate that, in addition to equity-based motivations, liquidity constraints are an important margin of the decision to default when default costs fall. Although default costs act as a wedge between negative equity and default, they also reduce a household's ability to smooth consumption. In addition to correcting foreclosure externalities, debt relief policies aim to help borrowers address problems arising from incomplete debt contracts and, in turn, smooth consumption in a crisis (Bolton and Rosenthal (2002), Piskorski and Seru (2018)).

[6] Fuster and Willen (2017) argue that a consumption-portfolio choice model with constraints can replicate such behavior. Because of high marginal utility of current consumption, liquidity-constrained borrowers discount the investment payoff from the option to a greater extent than their unconstrained counterparts with an identical contract and similar default costs, and therefore, act as if impatient.

For example, HAMP sought to both reduce foreclosures and stimulate consumption (Ganong and Noel (2020)). My estimates of the liquidity margin are consistent with a model in which borrowers with high marginal utility of consumption default to better smooth consumption. These borrowers are likely to have a high marginal propensity to consume out of reduced repossession risk.

Finally, my results demonstrate using quasi-experimental variation that a household's sensitivity to default is related to the costs of this decision. Estimates of the sensitivity of defaults to these costs are important to address questions about how households make decisions about financial leverage (Bailey et al. (2018)), which ultimately can play a role in the both booms and busts of house price bubbles (Geanakoplos (2010)).

The rest of this paper is organized as follows. Section I provides background on the natural experiment. Section II discusses the empirical research design and data used. Section III introduces the regression specification and presents the main results. Section IV discusses threats to identification and alternative modes of inference. Section V considers the magnitude of the effect and external validity, as well as heterogeneous treatment effects. Finally, Section VI concludes.

I. The Dunne Judgment and the Irish Mortgage Market

In early summer 2011, at the peak of the Irish financial crisis, several appeals against residential repossession proceedings were heard by the High Court in Dublin. Irish house prices were down 50% from their credit-boom peak, unemployment tripled to 14%, and mortgage defaults and foreclosures were climbing. On July 25, 2011, Ms. Justice Dunne ruled in the case of *Start Mortgages Ltd & Ors v Gunn & Ors* that a 2009 Government Act intended to update Irish repossession law had unintentionally repealed repossession law. The landmark ruling effectively removed any risk of repossession for mortgage borrowers in the short run, banks could no longer petition the Irish courts to enforce their rights to collateral repossession. Popular online media outlets reported on the judgment that night and all major national newspapers in Ireland carried the story the following morning.[7]

Borrowers were aided in the appeals cases by a group of lawyers and business people called *New Beginning*, which further helped publicize the ruling. For example, national broadcaster RTÉ stated in the afternoon of the ruling that "Hundreds of other people who are in arrears with mortgages created before December 2009 could also be affected by the decision. New Beginning says it puts a question mark over the entire repossession system since 2009." This third sentence demonstrates clearly the ramifications of the ruling for

[7] Gartland, Fiona, 2011a. Challenges delay repossession cases, *The Irish Times* 4, July 4. Gartland, Fiona, 2011b, Hundreds of home repossession cases may be struck out, *The Irish Times* 6, July 26. McDonald, Dearbhail, 2011, More than 100 home repossessions ordered by the courts could be legally challenged after a landmark ruling issued yesterday, *Irish Independent* 8, July 26. RTE, 2011, Ruling to put repossessions in doubt? July 25, 2011.

borrowers with mortgages originated before December 1, 2009, the cutoff date that I use for identification in this study.

The source of the legal lacuna begins with the Land and Conveyancing Law Reform Act 2009 (henceforth "Land Act"). The Land Act was passed by the Irish legislature in 2009 to replace older legislation used in repossession proceedings, namely, the Registration of Title Act of 1964. Section 62(7) of the 1964 law was replaced by article 8 in the 2009 Land Act. Nearly two years later, Justice Dunne noted in her ruling on the 2011 case that "those lenders who did not have an entitlement to apply for an order pursuant to s. 62(7) by December 1, 2009, *are not in a position to avail of the provisions of the 2009 Act to apply for an order of possession as their right to apply for such an order is not saved* by the provisions of the 2009 Act" (emphasis added). In other words, if a lender had not applied for possession of the collateral under section 62(7) of the old law by December 1, 2009, they could not do so after. However, loans issued after this date were still covered by the new 2009 law and lenders could still petition the court for repossession after Justice Dunne's decision in 2011.

Newspaper and online evidence from the time suggests that the cutoff in the ruling was salient. In addition to RTÉ, Fiona Gartland, writing in *The Irish Times* the following day, noted that "The new Act was introduced on December 1st, 2009, and it only applied to mortgages created after that date, the judge found."[8] Following up the next month, Gartland wrote about the immediate impact of the ruling and noted the *New Beginning* group's claims of an increase in borrower inquiries after the judgment.[9]

How was it that lenders continued to issue foreclosure notices against borrowers when the law was defunct? In response to the initial working paper of this study (O'Malley (2018)), a member of the *New Beginning* group, Andrew Robinson, wrote in the *Irish Independent* that lenders were in such haste to issue foreclosures in 2010 and 2011 that they failed to take adequate care in issuing the pleadings.[10] He claims that Irish banks' strategy for addressing the increase in defaults at the time was to issue tenders for high-volume, low-price work, with the resulting lack of attention to detail meaning that nobody noticed the legal lacuna. Robinson claims that, in effect, lenders "sought an illegal order when a cursory check of their pleadings would have uncovered the error."

Although it is hard to quantify how many foreclosures were avoided by the ruling, the flow of repossessions was halted. Data from the Irish courts show that as the stock of mortgages in default continued to head towards 150,000 (roughly one-in-eight outstanding mortgages), the number of orders *granted* for possession fell in 2011 and continued to fall thereafter while the loophole continued to exist (Arthur Cox (2016)).

[8] Gartland, Fiona, 2011b, Hundreds of home repossession cases may be struck out, *The Irish Times* 6, July 26.

[9] Gartland, Fiona, 2011c, Loophole prompts surge in calls to mortgage group, *The Irish Times* 4, August 5.

[10] Robinson, Andrew, 2018, Courts can hold heads high in repossession mess, *Irish Independent*, February 14.

Unlike the majority of their U.S. counterparts, Irish mortgages are recourse in nature, that is, lenders can pursue borrowers for any residual sums owed after they have taken possession of the collateral. This fact, together with the uncertainty surrounding the length of the repossession moratoria means that a borrower was not necessarily free from repossession risk. It seems likely that borrowers understood that the ruling would be temporary, with legislation enacted to correct it. Commenting on the latter prospect, O'Neill (2011) writes that "There have already been some indications of possible appeals of Ms Justice Dunne's decision. However, such an appeal will not occur quickly and, in any event, the decision may very well be upheld – leaving aside its dramatic consequences, prima facie, it appears well reasoned." The loophole was closed two years later, when in July 2013, a new Land and Conveyancing Law Reform Act was signed into law by the legislature. After this law went into effect, the number of repossession cases in the Circuit Court nearly quadrupled from 2013 to 2014 (Arthur Cox (2016)).

II. Data and Sample Selection

The data comprise a panel data set of loans from the Central Bank of Ireland's loan-level database, which were collected for stress testing the major Irish lending institutions during the recent financial crisis. The data cover four large banks, which together cover about two-thirds of the Irish mortgage market at the time. The data are observed quarterly from one year prejudgment to one year post, and thus span eight quarters in total from October 2010 to July 2012. In each quarter, the data contain up to date information on loan performance, estimated LTV ratios, interest rates, and product types. The data also contain detailed information recorded at loan origination, including the origination date and details about the borrower, their total income, their NUTS3 region of residence, and their year of birth.[11]

As part of the research design, I limit the sample to loans originated 180 days either side of the cutoff date specified in the Dunne judgment. I also limit the sample to mortgages taken out by homeowners rather than investors. All loans in the sample are primary loans, meaning that they have no secondary loans attached to them, nor are they equity release loans from a borrower's larger portfolio, which made up roughly one-third of loan originations during the period (Irish Banking Federation (2010)). If a control-group loan were originated for equity release or as part of a larger loan portfolio, then the effect of the Dunne judgment on the borrower's incentives would be less clear.

Since the assignment of loans to treatment and controls groups is not truly random, I use a matching algorithm to make the two groups observably similar at the first observation. Though the identification is relatively clean, there nonetheless remains the possibility that the two groups are not comparable on average pretreatment. To mitigate this concern, I apply a two-step greedy

[11] Nomenclature des Unités Territoriales Statistiques (NUTS) is the European Union geocoding standard. There are eight NUTS3 regions in Ireland, of which Dublin Region is the largest.

matching procedure, first estimating a treatment-propensity score as a logistic function of observable factors associated with loan default, and then dropping observations that do not receive a nearest-neighbor match in the second stage (Imbens and Rubin (2015)). Together with regression adjustment in Section III, this matching procedure ensures that the estimates are doubly robust (Imbens (2004)).[12] The results without matching are qualitatively similar.

The data set used in the remainder of this paper consists of 80,272 loan-time observations. Of the 7,913 sample loans, 4,488 are no longer subject to repossession after the Dunne judgment.

Table I reports both the overall mean and the group means for relevant loan characteristics. As the loan sample is defined narrowly and matched, there are few meaningful differences and most are mechanical and due to the treatment group loans being older by definition. For example, treatment loans are associated with lower balances, older borrowers, and lower terms to maturity. Borrowers in the treatment group do face slightly higher interest rates on average, but this translates into only small differences in their average monthly installment.

Table I also reports the proportions of discrete variables by treatment and control groups. Treatment and control observations are generally balanced across regions, but control observations are more likely to be first-time home purchasers. Considerable imbalance exists between fixed and variable interest rates across the two groups, however, which if not statistically accounted for would be a valid threat to identification.

III. The Impact of the Dunne Judgment on Mortgage Defaults

To test whether the Dunne judgment had a causal impact on mortgage defaults, I model the default (flow) outcomes of the two groups over the pre- and postjudgment periods. The outcome of interest is the default transition of loan i in treatment group j at time t,

$$\text{Default}_{ijt} = (90\text{DPD}_{ijt} = 1 \,|\, 90\text{DPD}_{ijs} = 0 \quad \text{for all } s < t).$$

At each point in time, the dependent variable takes the value of 1 if the loan *transitioned* to nonperforming status since the last observation, where nonperforming status is defined following Basel as 90 days past due. I define the transition to default as an absorbing state. Thus, once a loan is coded as in default, it drops out of the sample thereafter.

[12] The factors used in matching are issuing bank, interest rate, interest rate type (fixed, variable), original LTV ratio, income, and NUTS region. I use the R package *MatchIt* to perform the matching (Ho et al. (2011)). Greedy matching entails matching observations without replacement. The procedure matches the first treatment observation to a control and then continues to the next unmatched treatment observation, ignoring the effect of the previous step on potential subsequent matches. As a result, the algorithm does not attempt to optimally match observations between groups, but it is easier to implement and the loss over optimal matching is usually small (Imbens and Rubin (2015)).

Table I
Summary Statistics

This table reports mean values for continuous control variables (top panel) and proportions in percent for discrete control variables (bottom panel). Estimates are computed on data in the period immediately preceding the judgment (March 2011). Also reported is the variable dictating treatment status, that is, distance from cutoff, which is given as the number of days between the loan origination date and the December 1 cutoff date.

	Mean		
	All	Control	Treatment
Treated	0.50		
Interest rate (%)	4.11	3.89	4.33
Loan-to-value ratio (%)	77.42	78.59	76.25
log(Income)	10.95	10.96	10.95
Borrower 1 year-of-birth	1973	1973	1972
Outstanding balance (€)	176,080.10	180,299.20	171,861.10
Term remaining (months)	276.26	283.38	270.12
Monthly installment (€)	952.54	939.62	965.49
Distance from cutoff (Origination days since 1st December 2009)	−4.12	88.9	−97.2

	Proportion (%)		
	All	Control	Treatment
Fixed interest rate	40.34	55.56	25.11
First-time purchaser	49.57	53.06	46.07
NUTS3 region			
Border	10.64	9.56	11.71
Dublin	30.01	32.81	27.22
Mid East	13.64	13.49	13.79
Mid West	9.65	9.12	10.18
Midlands	5.03	4.71	5.36
South East	8.76	8.09	9.43
South West	13.73	13.76	13.71
West	8.54	8.48	8.61

Figure 1 plots the empirical default rates by group over the sample period. While they follow parallel trends before the judgment, the default rates then diverge after the judgment. The treatment group's default rate rises by roughly half a percentage point in the first period after the ruling, consistent with a causal impact of removing repossession risk on mortgage default. The default rate for the treated loans remains elevated over the remainder of the sample period.

A. Regression Model

To test whether the judgment had a statistically significant impact on defaults and to adjust for potentially confounding variation, I estimate several

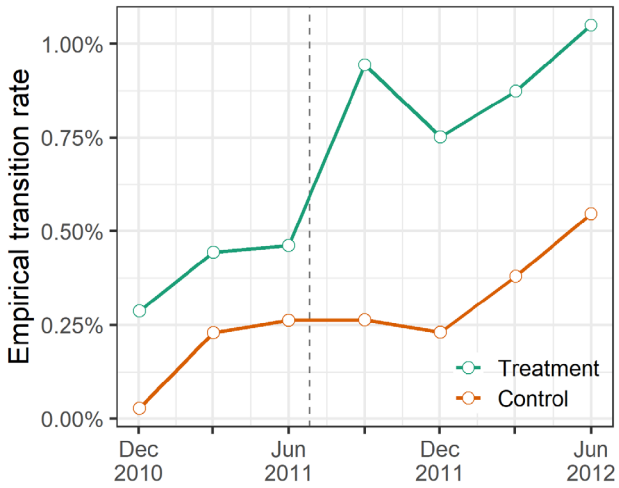

Figure 1. Empirical default rates by group. This figure plots the empirical default rates over time by treatment group. Data are quarterly transition-to-delinquency rates. The dashed vertical line indicates the date of Dunne judgment ruling. The treatment groups comprise loans originated up to six months before the cutoff date specified in the judgment. The control group comprises loans originated within the six months after the judgment. (Color figure can be viewed at wileyonlinelibrary.com)

regression models. The models are variations of the panel DD specification

$$\text{Default}_{ijbgrt} = \alpha + c_i + \beta^{DD}(\text{Treatment}_j \times \text{Post}_t) + \mathbf{X}'_{it}\Psi + \phi_{rt} + \tau_{bgt} + \epsilon_{ijbgrt}, \quad (1)$$

where i, j, b, g, r, and t are index loan, treatment group, issuing bank, product type, borrower region, and time, respectively, and β^{DD} is the treatment effect of interest to be estimated from the data. Under the assumption that the error term is exogenous conditional on the model covariates, $\beta^{\hat{DD}}$ is the causal effect of the Dunne judgment on default and hence an estimate of the sensitivity of mortgage default to the removal of repossession risk.

The specification assumes that for each i, j, g, b, r, and t, the probability of mortgage default is an additive function of a loan effect that controls for the time-invariant risk of each loan, a matrix of loan characteristics that are associated with loan default (e.g., negative equity), group-time fixed effects, an idiosyncratic loan-level error term, and an additive treatment effect for the impact of the Dunne judgment (β^{DD}). Identification of β^{DD} relies on a comparison of the within-loan variation in default outcomes over time between the treatment and control groups. This specification also assumes no effect of the origination date on loan default beyond the effect absorbed by a treatment group or loan dummy, which is equivalent to assuming that loans issued at different points in time have a constant default probability differential, an assumption that I relax below by adjusting for a linear time trend interacted

with month-of-origination dummies. In that case, identifying variation comes from deviations from this trend in the postperiod between treatment and control loans.

Table I shows that the group compositions are generally imbalanced across product types, which may present a threat to valid causal inference. For example, the 2011 European Central Bank (ECB) interest rate hikes may be more likely to be passed through to the treatment group, which would confound the estimated treatment effect. In the Irish market, because banks retain the option to change interest rates on variable rate loans rather than maintain a fixed margin, pass-through is not complete. I therefore also estimate specifications that contain a product type $\times$ issuing bank $\times$ time fixed effect. This term adjusts nonparametrically for any confounding variation coming from changes in bank-specific interest rate policies by isolating variation at the issuing bank-product-type level. Intuitively, by comparing default rates for variable-rate mortgages in Bank A at time T across treatment and control groups, this variation should not bias the estimate. In several specifications, I also adjust for the level of the interest rate through the main effect or the borrower's payment amount, though there is little remaining variation in specifications that feature this high-dimensional fixed effect. To account for within-loan correlation in the error term, I cluster all specifications at the individual loan level.

B. Regression Results

Table II presents results of variations of equation (1) estimated by ordinary least squares. Column (1) reports results from a model with only the basic ingredients of the DD model: an intercept, a treatment group dummy, a postperiod dummy, and the interaction between the two. In each subsequent column, I add or subtract various combinations of covariates and fixed effects to test the robustness of the estimates.

Starting with specification (6), I allow the baseline transition rate to vary with vintage by first adding origination-month fixed effects. In specifications (7) and (8), I add vintage effects such as a linear trend for each monthly cohort. Specification (9) presents the results after including high-dimensional issuing bank $\times$ product type $\times$ time fixed effects. Specifications (10) and (11) add a linear interaction between origination date and calendar time, (12) adds vintage month dummies, (13) adds a linear origination date term interacted with the post dummy, and the final specification contains a linear interaction between vintage and calendar time.

Across all specifications, the DD coefficient is shown in the first row and is statistically different from zero at the 5% level in all specifications bar the final two, which are significant at the 10% level. The estimated treatment effect is roughly 0.3 percentage points for most specifications. The estimates of 0.5, 0.43, and 0.48 percentage points in specifications (8), (11), and (12), respectively, are higher than the rest. This evidence suggests that the Dunne judgment had a causal impact on the estimated counterfactual default rate of approximately 0.3 percentage points. Using the parameters of specification (1)

Table II

Difference-in-Difference Regression Results

Each column reports results from a variation of equation (1). The bottom rows of each panel isolate which fixed effects are included in each specification. If a cohort trend is included in the model, the bottom row also specifies the type. The coefficient on Treated × Post, (β^{DD}), is the coefficient of interest and is isolated in the first row of both panes. The remaining estimates are coefficients on control variables. Loan-clustered standard errors are shown in parentheses. Significance levels 10%, 5%, and 1% are denoted by *, **, and ***, respectively.

	Default						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treated×Post	0.0034***	0.0036***	0.0038***	0.0038***	0.0033***	0.0034***	0.0033***
	(0.0009)	(0.0009)	(0.0009)	(0.0009)	(0.0009)	(0.0009)	(0.0009)
Treated	0.0021***	0.0007	0.0006	0.0006	0.0022***		
	(0.0006)	(0.0006)	(0.0007)	(0.0006)	(0.0006)		
Post	0.0016***	−0.0001	0.0004	0.0005		0.0016***	
	(0.0005)	(0.0005)	(0.0005)	(0.0005)		(0.0005)	
Installment (€)		−0.000001					
		(0.000001)					
Loan-to-value (%)		0.0001***					
		(0.00001)					
Negative equity		0.0007	0.0037***	0.0037***			
		(0.0008)	(0.0007)	(0.0007)			
log(Outstanding balance)			0.0042***	0.0041***			
			(0.0005)	(0.0005)			
Borrower 1 year-of-birth		−0.0002***	−0.0002***	−0.0002***			
		(0.00004)	(0.00004)	(0.00004)			
log(Income at origination)		−0.00003	−0.0024***	−0.0024***			
		(0.0007)	(0.0006)	(0.0006)			
Interest rate			−0.0008	−0.0011**			
			(0.0007)	(0.0005)			
Above median interest rate			0.0011				
			(0.0008)				
Variable rate		0.0038***	0.0038***	−0.0053*			
		(0.0005)	(0.0005)	(0.0027)			
Variable rate×interest rate				0.0023***			
				(0.0007)			
Fixed effects							
Region		Y	Y	Y			
Time					Y		
Vintage (month)						Y	Trend
Observations	80,259	80,084	80,251	80,251	80,255	80,259	

(Continued)

Table II
(Continued)

	(8)	(9)	(10)	(11)	(12)	(13)	(14)
Treated×Post	0.0050**	0.0029***	0.0035**	0.0043***	0.0048**	0.0035*	0.0035*
	(0.0020)	(0.0009)	(0.0014)	(0.0015)	(0.0020)	(0.0019)	(0.0019)
Treated	0.1149	0.0013**	−0.0001				
	(0.1298)	(0.0006)	(0.0012)				
Installment (€)						−0.00003***	−0.00003***
						(0.00001)	(0.00001)
Loan-to-value (%)						0.0007***	0.0007***
						(0.0001)	(0.0001)
Negative equity						0.0021	0.0022
						(0.0014)	(0.0014)
log(Outstanding balance)			0.0042***	0.0041***			
			(0.0005)	(0.0005)			
Interest rate			−0.0008	−0.0011**			
			(0.0007)	(0.0005)			
Above median interest rate			0.0011				
			(0.0008)				
Variable rate						0.0041	0.0044
						(0.0031)	(0.0032)
Origination date			−0.0001				
			(0.0003)				
Origination date×time			0.0000	−0.000000**			
			(0.000000)	(0.000000)			
Origination date×post						−0.00002**	
						(0.00001)	
Fixed effects							
Loan				Y	Y	Y	Y
Time	Y	-	Y	Y	Y	Y	-
Bank-Product-Time		Y					
Vintage (month)	Trend				Y		
Vintage (Linear)-Time							Y
Observations	80,259	80,259	80,259	80,259	80,259	80,084	80,084

to compute the counterfactual in the postperiod, this result translates into a relative effect of around 60%.

The initial gap in the default rates between treatment groups in Figure 1 is a concern, since we would not expect to see such a difference between loans issued so close together. The results in Table II suggest that either these loans were observably worse on origination characteristics or they were worse credit risk on unobservable characteristics. For example, the treatment parameter, which measures the preperiod average difference in default rates, is statistically greater than zero in the baseline specification but becomes zero after covariates are adjusted for or the baseline hazard is allowed to vary with vintage effects. In both cases, the estimated DD parameter is stable and statistically significant. The remaining specifications argue against this gap biasing the treatment effect. If the treatment group diverged in the counterfactual scenario regardless of treatment, due to unobserved differences, we would expect the vintage trends to absorb this. For example, in specifications with cohort effects, the treatment effect is identified from deviations in the preperiod trend. But there is no reason why deviations from trend should occur for only the treatment group.

C. Event-Study Specification

In the previous section, I describe how the DD specification computes the counterfactual outcome of the treatment group using the posttreatment outcome of the control group and the assumptions needed to identify the causal impact of the judgment on default. However, a natural analysis of the research design is to compare default rates between treatment and control groups at each point in time, rather than between the pooled pre- and postjudgment periods. The key implication of the hypothesis of this paper is that default rates should diverge only after the judgment in July 2011. An event-study model differs from equation (1) in that the $\text{Treatment}_g \times \text{Post}_t$ parameter is replaced by a full set of treatment by observation quarter interactions. The resulting estimator allows free coefficients for each time period's treatment effect: periods before the judgment should show no difference in outcomes from the first-period difference between treatment and control groups. Though this estimator will be less efficient than that in equation (1), it provides a test of the parallel trends assumption of the standard model. Of course, parallel trends in the postperiod cannot be tested, but the model here can provide evidence.

Specifically, I estimate the following specification:

$$\begin{aligned}\text{Default}_{ibfgrt} = \alpha + c_i + \sum_{t \neq \text{Dec2010}} \big\{ & \beta_t \times \mathbb{1}(\text{Quarter}_t) + \beta_t^{DD} \\ & \times \mathbb{1}(\text{Quarter}_t) \times (\text{Treated}_g) \big\} + \mathbf{X}_{it}' \Psi + \phi_{r,t} + \tau_{b,f,t} + \epsilon_{ibfgrt}. \end{aligned} \tag{2}$$

The hypothesis here is that for all quarters before July 2011, β_t^{DD} should be statistically zero and some subset of the β_t^{DD} coefficients in the postperiod

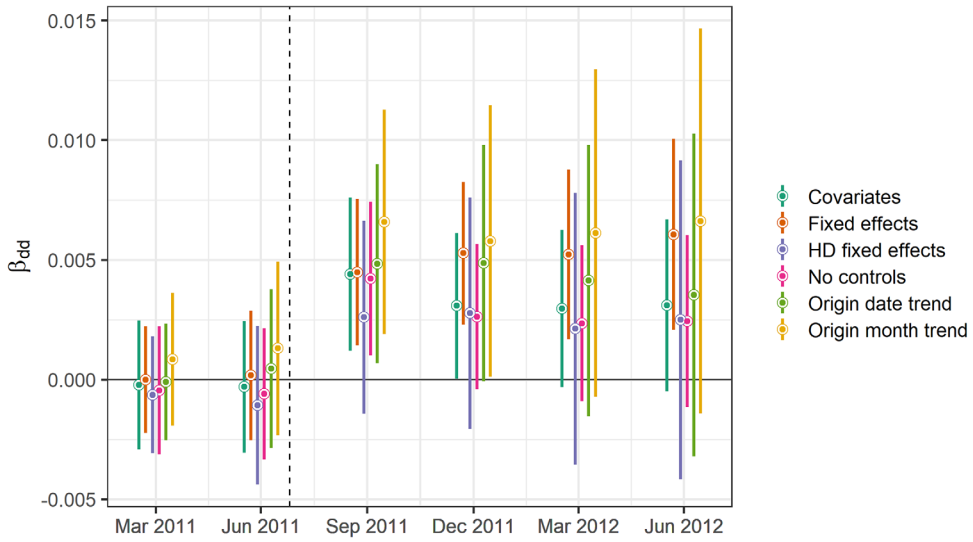

Figure 2. Event-study coefficients. This figure shows the coefficients on the interaction between calendar-quarter fixed effects and the treatment dummy. In each period, the point estimates come from six different specifications, labeled in the legend. The 95% confidence intervals calculated using loan-clustered standard errors are also depicted. (Color figure can be viewed at wileyonlinelibrary.com)

should be statistically greater than zero. The estimated counterfactual for the treatment group is quite flexible in this model. The first period β^{DD} is constrained to be zero by construction, and in all subsequent periods, the change in the control group's time trend is used to estimate the counterfactual change in the treatment group. Any deviation from this prediction is considered the causal impact of the judgment. This is a strict test—any significant deviation from the prediction in the two periods before the judgment will show up as a nonzero treatment effect.

The identifying assumptions from Section II.A continue to hold here: any group-time varying confounders must be accounted for. For this reason, this model inherits the same set of control variables and fixed effects as equation (1). Inference remains straightforward and standard errors are clustered at the individual loan level. The sources of variation in this model are similar to those in the previous pooled model: the treatment effect is identified from the within-loan default variation away from its sample mean between the treatment and control groups at each point in time.

The results of several event-study specifications are depicted in Figure 2. Each date on the x-axis provides the estimate of the nonparametric treatment effect, with the vertical dashed line highlighting the Dunne judgment. Each β_{DD} shows the difference in default rates from the first-period difference. The figure clearly shows that the treatment groups followed parallel pretrends that diverge in the first period after the ruling. Across the six specifications

considered, the first postperiod effect is generally the most precisely estimated with confidence intervals widening afterward. However, the pattern is similar across specifications, with the treatment effect occurring only after the time of the ruling.

D. Month-of-Origination Nonparametric Treatment Effect

To further rule out competing explanations for the treatment effect, similar to the event study, I now allow the treatment effect to vary nonparametrically in the month of origination instead of the quarter of observation.[13] This test should reveal a significant treatment effect only for the months prior to December 2009, the cutoff month. Identification in this model is similar to that in a regression discontinuity design (RDD) model in that the treatment effect is expected to jump discretely in the first month before the cutoff. Unlike the standard RDD, however, this hybrid model identifies the effect from the cross-sectional *difference* in the relationship compared to the December effect, after the event.

Specifically, I estimate the following specification:

$$\text{Default}_{iot} = \alpha_i + \sum_{o \neq \text{Dec 2009}} \{\beta_o \times \mathbb{1}(\text{Origination}_o) + \beta_0^{DD} \times \mathbb{1}(\text{Origination}_o) \times \text{Post}_t)\} + \epsilon_{iot}. \quad (3)$$

Estimates of the β_0^{DD} coefficients for four specifications are shown in Figure 3 along with the equivalent pooled effect for each specification. The patterns are generally consistent with the treatment effect being present only in treatment months, although the effect is stronger in earlier months. It is also possible that earlier months have higher treatment effects due to unobserved borrower differences. Therefore, in further specifications, I add vintage × calendar time parameters to control for differential pretrends for each cohort. In this case, identification of the month-dummy interactions relies on deviations from their prejudgment linear trend. In Panel A, in which there is no attempt to control for pretrends, earlier months have statistically larger treatment effects than later months. However, this difference reverts to the mean in Panel B, in which I adjust for cohort trends. The individual month dummies are rarely statistically significant at conventional levels, although the pooled effects are significant in each case.

An equivalent regression discontinuity estimate shows no statistical effect of the treatment. This could be due to a lack of power in the regression discontinuity estimate. In a Monte Carlo simulation, I find that specifications similar to equation (3) have similar mean-square error and bias, compared to appropriate regression discontinuity estimates. Therefore, the weak treatment effects observed in Figure 3 for November and October could be why the regression

[13] I thank the referees for suggesting variations of this test and the discussion in this section.

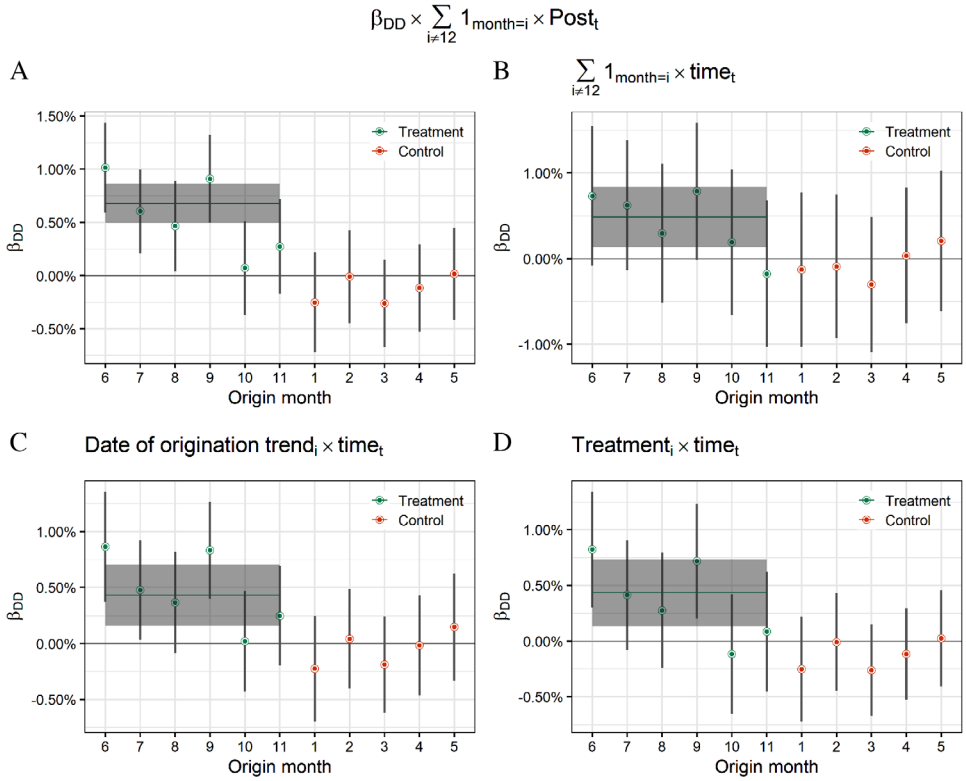

Figure 3. Month-of-origination nonparametric treatment effects. This figure plots the month-of-origination nonparametric treatment effects. Each point represents a coefficient on the effect of each month, relative to December 2009 (12 is omitted). The estimates come from regressions based on equation (3). The β_0^{DD} coefficients, which correspond to origination-month treatment effects, are shown for different specifications, as are pooled estimates (horizontal lines). Each model controls for cohort-time effects. The functional form is shown in the caption of each panel. (Color figure can be viewed at wileyonlinelibrary.com)

discontinuity estimate finds no effect. Since regression discontinuity assumes that the effect varies linearly with loan origination month, it will show a null treatment effect for the patterns observed in Figure 3.

IV. Robustness of Treatment Effect

To gauge the sensitivity of the treatment effect, I perform two additional analyses. In the first, I change the bandwidth of the sample selection procedure and allow more or less loans to enter the analysis. I then examine how the estimated effect changes. In the second analysis, I conduct a permutation exercise to test the null hypothesis of no effect of the Dunne judgment on mortgage defaults.

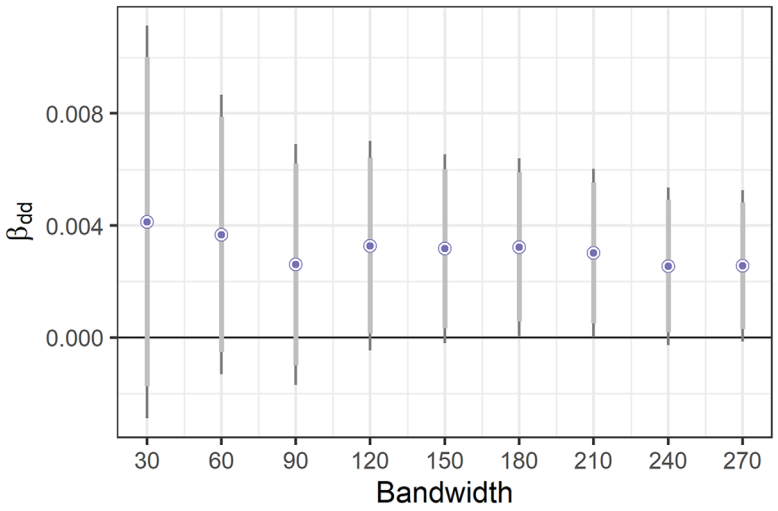

Figure 4. Sensitivity of the treatment effect to chosen bandwidth. This figure shows how the treatment effect changes in response to changes in the estimation bandwidth. Each point represents the estimated treatment effect from a model estimated on a new sample using a different bandwidth. The bandwidths are shown on the *x*-axis. The model is similar to that in equation (1), but includes a cohort-specific linear time trend. Vertical lines indicate 90% (light gray) and 95% (dark gray) confidence intervals. (Color figure can be viewed at wileyonlinelibrary.com)

A. Varying the Bandwidth of Estimation

A reasonable specification should not permit large changes in the treatment effect when the parameters change. One obvious parameter to change is the bandwidth used in the sample selection procedure, as described in Section II. Since the threshold of 180 days is essentially arbitrary, I repeat the analysis on bandwidths from 30 to 270 days, in 30-day increments. I rerun the model in specification (1) after allowing each origination-month cohort to have its own pretreatment trend. This seems sensible as it is unlikely that vintage cohorts will be on parallel default trends as the bandwidth is progressively widened.

Results, as depicted in Figure 4, show that the treatment effect is stable across specifications. The treatment effect is not statistically different from zero in the smaller bandwidths, possibly due to small sample size.

B. Permutation Inference

Since treatment is applied at the group level, there are essentially two independent clusters under study, which complicates conventional statistical inference. One alternative mode of inference is randomization or permutation-based inference. In a randomized experiment, permutation inference is undertaken by repeatedly randomizing the treatment assignment vector, reestimating and storing the treatment effect under these different permutations, and

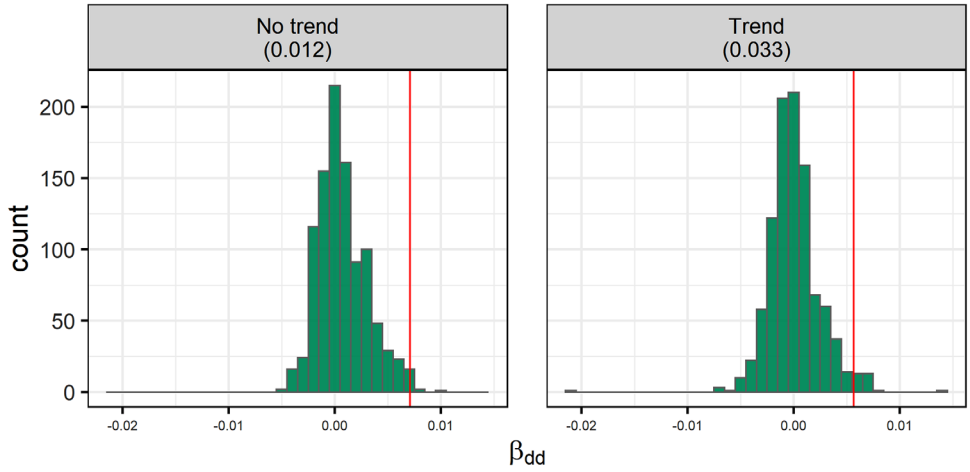

Figure 5. Permutation exercise. This figure shows two null distributions of 999 "placebo" analyses. The left panel shows the analysis for a model with no cohort trends; the right shows the same model with origination-month trends. True treatment effects are indicated by horizontal red lines. Associated *p*-values are provided below the plot titles in parentheses. (Color figure can be viewed at wileyonlinelibrary.com)

calculating the *p*-value as the location of the true effect in this simulated null distribution.

I conduct a similar permutation exercise to estimate the significance of the effect size by simulating a null distribution using many "fake" natural experiments. Specifically, I randomly take 999 unique loan origination dates (sampled without replacement) from the database of all outstanding loans in Ireland during the sample period and analyze the corresponding placebo treatment effect. The idea is to test whether my research design often demonstrates positive treatment effects in the Irish mortgage data, even when there should be no effect at all. Inference in this case is similar to Ganong and Jäger (2018), who provide theoretical and simulation evidence on the effectiveness of randomization inference in the analysis of regression kink designs.

To create the estimates for each fake experiment, I follow the same procedure as in the main analysis: I separate the data into treatment and control cohorts 180 days either side of the cutoff, I select the sample using matching, I estimate the DD regression model and store the estimates, and I calculate the *p*-value by measuring the position of the absolute value of the true estimate under the null distribution. The *p*-value is therefore the probability that I would observe at least as large an effect, when there is, in fact, no effect.

Figure 5 presents the results for two models, one with and the other without an origination-month time trend. Both histograms of null distributions show that the observed effects, with *p*-values below 5%, are unlikely to have occurred by chance.

V. Magnitude and Mechanisms

A. Discussion of Mortgage Default Theory

In double-trigger models, borrowers generally default when the option value of repayment net of the interest payment falls below their own personal default costs (e.g., Foote and Willen (2018), Ganong and Noel (2020)). In other words, for these borrowers, the discounted expected financial gain from paying the mortgage no longer justifies repayment, compared to their outside option of default.[14] For example, a borrower whose mortgage is deep underwater will default if rent is cheap compared to their interest payment, whereas a less underwater borrower will not default if they believe that they will benefit from future house price appreciation by continuing to pay over the odds for the house. Similarly, myopic borrowers who discount the option value of repayment heavily may default if interest payments are high compared to default costs.

As the above examples suggest, two types of borrower are on the margin of defaulting at current default costs: borrowers in deep negative equity who expect little financial gain from repayment, and liquidity-constrained borrowers who discount the future to a high degree. When a borrower is in deep negative equity, the mortgage is less likely to generate a positive investment profit. In this case, the repayment option is deep out-of-the-money and the borrower is better off defaulting in a purely financial sense.[15] Default costs lead some of these borrowers to continue to pay. Therefore, when default costs fall, borrowers on this margin default.

Liquidity-constrained borrowers, in contrast, generally do not have the ability to choose when to default. These borrowers default because they receive an income shock and they simply cannot afford to make the required payments. However, some liquidity-constrained borrowers will be able to choose when to default. These latter borrowers continue to repay by cutting consumption to avoid the costs of default. Because they discount the expected future payoff from repayment so highly, these borrowers would default if not for the default costs.[16] When these costs drop below some threshold, repayment by sacrificing current consumption is less valuable than the default option.

The policy implications of the excess default that I document in Section III are affected by which margin contributes more strongly. Depending on the benefit of smaller foreclosure externalities, the moral hazard cost may be large

[14] Negative equity is a necessary condition for default in these models, since positive equity would allow a borrower to avoid default costs by selling the home.

[15] Since repayment today allows a borrower to choose whether to repay or default next period, mortgage contracts contain some optionality. Though mortgage contracts in Ireland contain recourse, default costs are limited by the borrower's downside in bankruptcy. Therefore, repayment is an asymmetric bet on future house prices. If prices remain stable or fall, then the borrower is no worse off. However, if prices rise, the borrower profits if the increase is sufficient to lead to positive equity. This asymmetry resembles the payoff from a call option. See Foote and Willen (2018) for an extensive discussion of the option-theoretic approach to mortgage valuation.

[16] In a consumption-based portfolio choice model with constraints, liquidity shocks lead borrowers to have high-current marginal utility of income and therefore discount the future highly (Fuster and Willen (2017)).

enough to dissuade policymakers from implementing a foreclosure moratoria. There is some subtlety to this cost though. If the policy only induces liquidity-constrained borrowers to default, then these borrowers are likely to have a high marginal utility of current consumption and therefore a high marginal propensity to consume. By allowing these households to share their income risk more cheaply, the policy may induce additional consumption.

In this section, I provide evidence on the two margins and discuss interpretation of the estimated magnitude from Section III.

B. Magnitude of Effect and External Validity

As outlined above, double-trigger models suggest that borrower defaults following a decrease in default costs are due to both equity and liquidity-based concerns. The distributions of borrower equity and available liquidity thus affect the external validity of this paper. Implementing a foreclosure moratorium in a country where the labor market is strong and borrowers have substantial equity in their homes would likely not increase the delinquency rate in any economically meaningful way. In 2011, the Irish economy had recently suffered a collapse in real estate prices (homes lost approximately half of their value over the 2008 to 2012 period) and a tripling of the unemployment rate. Both of these factors increased the likelihood of excess delinquency following the removal of repossession risk. Although different countries may not have observed an effect as large as that documented here, such a macroeconomic crisis is exactly when policymakers would wish to implement a debt relief policy. This episode thus represents an ideal setting to study the effects of reducing default costs, such as repossession risk.

The external validity of the result above also relies heavily on the legal and societal context in which it was observed. These factors influence how borrowers perceive the remaining costs of default once repossession risk is removed. Mortgage contracts in Ireland are recourse in nature, and therefore, incentives to strategically default are muted because a borrower cannot "walk away" from residual debt balances. Although Irish borrowers have strong legal protections in debt proceedings, lenders can access a borrower's assets through an *execution against goods* order (Citizens Information (2016)). If borrowers believed that the Dunne judgment period was likely to be temporary, they would have expected to observe these costs at a later date. Absent rare garnishee orders, however, lenders do not have access to a borrower's future income. Borrowers facing deep negative equity or severe liquidity constraints might then expect their future default costs to be insignificant compared with their current situation.

If repossession risk were permanently removed with certainty, some underwater or constrained borrowers would still not default because of the negative effect on their future access to credit. This concern is particularly relevant for U.S. borrowers, where access to credit markets and loan terms is dependent largely on a borrower's credit score (Agarwal et al. (2017)). Creditworthiness factors similarly into credit supply decisions in Ireland. The Irish Credit

Bureau is an industry-run private firm that collects and supplies credit histories to member institutions.[17] Credit reports typically carry 24 months of repayment history and loans remain in the database for five years after the loan has expired regardless of whether the debt was fully repaid (Citizens Information (2018)). Such information sharing between lenders would further reduce the return on strategic default.

Moral and social costs are also a large impediment to mortgage default. In their wide-ranging study on strategic default, Guiso, Sapienza, and Zingales (2013) document that moral considerations influence the decision to default. Interestingly, they also document that perceptions about fairness temper the reluctance to default. Consistent with this evidence, during the height of the Irish crisis, Prof. Morgan Kelly argued in *The Irish Times* that not only would falling house prices contribute to an increase in the number of defaults, but also the growing perception that Irish banks' loose lending standards had caused the crisis in the first place.[18] Together with the rising number of borrowers defaulting before the Dunne judgment, such feelings likely made borrowers in Ireland more willing to default during the financial crisis.

My estimates of the average treatment effect of the judgment are smaller in absolute magnitude than the point estimate in Mayer et al. (2014), who find a 1.5 percentage point increase in the quarterly delinquency rate following the announcement of a mortgage modification program. The effect that they document translates into a 10% increase over the counterfactual, which is smaller than the 40% to 60% relative effect that I estimate here. Ghent and Kudlyak (2011) find a 30% relative increase in the delinquency rate in nonrecourse U.S. states compared to recourse states. To gauge the size of the effect in my sample, using control group data in a simple linear model, I estimate the association between default and LTV plus interest rates.[19] This cross-sectional comparison indicates that 0.3 percentage points are roughly equivalent to increasing the median LTV by 38 percentage points and the median interest rate by 2 percentage points.

Other factors that may contribute to the large effect that I document may arise mechanically. On average, the loans that I analyze are just over 1.5 years old at the time of the judgment. It is possible that low attachment to the home works in favor of a large treatment effect in this environment, consistent with the evidence in Guiso, Sapienza, and Zingales (2013). Since the loans that I analyze were issued at the end of the Irish credit boom when prices had already fallen, they are less likely to have substantial negative equity. This could favor finding a large effect if borrowers with extremely high negative equity default, even in the presence of high default costs—since these borrowers defaulted at

[17] The Central Bank of Ireland operates the more comprehensive Central Credit Register since 2017.

[18] Kelly, Morgan, 2010, If you thought the bank bailout was bad, wait until the mortgage defaults hit home, *The Irish Times* 11, November 8.

[19] I adjust for time-region fixed effects and also include a negative equity dummy to account for nonlinearity in LTV.

high rates before costs were reduced, they would not demonstrate a large relative treatment effect following the Dunne judgment. Similarly, if the sample contained a larger number of borrowers facing severe liquidity constraints, the treatment effect may be smaller because these borrowers default regardless of default costs.

C. Evidence on the Margins of Excess Default

C.1. Household Deposit Data

In addition to collecting data for stress testing the assets of the Irish banking sector, the Central Bank of Ireland collected data on the liabilities of the largest banks in December 2011. These data include account-level information on the deposits of Irish households at these banks. It is possible to link a subsample of borrowers in the loan data to these deposit data to obtain a measure of a household's liquid wealth on December 31, 2011, five months after the Dunne judgment. The data contain information on the average account balance over the 12 months from December 2010 to 2011, as well as the six months from July 2011 to December 2011. I use these data to calculate the implied average account balance in the first six months of the year, the six months roughly prior to the Dunne judgment. To aggregate up to the borrower level, I calculate the average across accounts, weighted by the account balances in December 2011. This variable measures household liquidity in the six months prior to the Dunne judgment. I then match this variable to the loan data in Section II by the unique borrower identifier. Figure 6 shows that the liquidity measure is distributed similarly across treatment groups.

There are some important caveats to these data. The first is that due to differences in reporting across banks, this matching is only possible for one bank in the sample, and therefore, any inferences are limited in their external validity. The analysis of heterogeneous treatment effects is also limited to a relatively small sample size of 1,347 loans, which means that estimates are likely to be noisy. The third caveat is that these estimates are static and do not take changes in liquidity into account. For example, a zero account balance may be normal for one borrower because they do not use this account, while it might signify binding liquidity constraints for another. Finally, these account balances could be endogenous if borrowers reduce their account balances in anticipation of default.

C.2. Heterogeneous Treatment Effects

I estimate a series of triple-difference regression models to separately estimate the magnitudes of the equity-based and liquidity-based reasons for the excess default. I first interact the DD variable with the terciles of the liquid

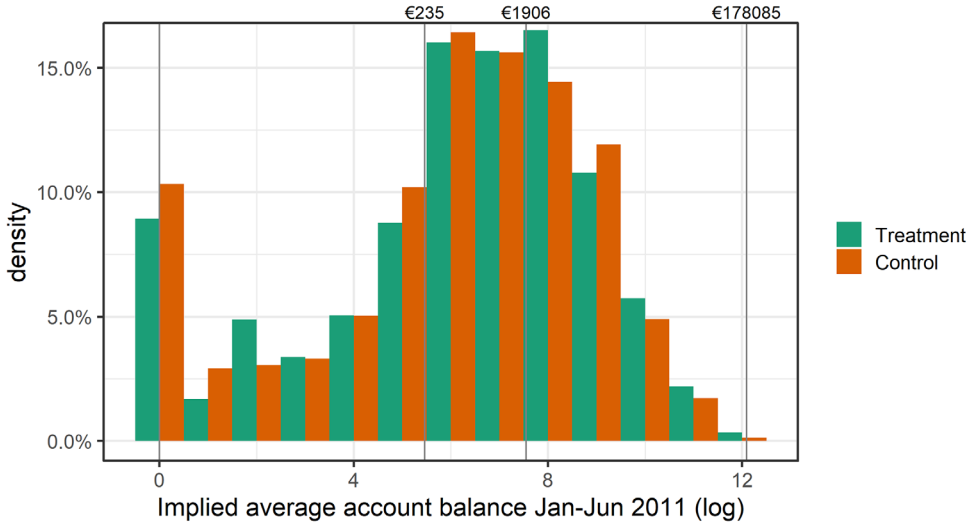

Figure 6. Distribution of borrower account data. This figure shows the estimated liquid account balances for borrowers, by treatment group. Data are the implied average account balance over January to June 2011, the six months before the Dunne judgment, and are calculated using the 12-month and six-month average balances in December 2011. $\text{Balance}_{Jan-Jun} = 2 \times \text{Balance}_{Jan-Dec} - \text{Balance}_{Jun-Dec}$. One is added to account balances before the natural log is calculated. Boundary values (in Euro) of the terciles are shown at the top of the grid. (Color figure can be viewed at wileyonlinelibrary.com)

wealth distribution.

$$\text{Default}_{ijt} = \alpha_i + \beta(\text{Treatment}_j \times \text{Post}_t) + \sum_{k \neq 1} \theta_k \times (\text{Treatment}_j \times \text{Post}_t \times \mathbb{1}_{T_{ik}}) + \phi_t + \epsilon_{ijt}, \quad (4)$$

where T_{ik} is a dummy variable that indicates whether borrower i has liquid wealth in the kth tercile of the distribution, and ϕ_t is a vector of time fixed effects. If the treatment effect is driven by a liquidity effect, then the coefficient on the DD term (β) should be positive and large (the effect when liquidity is in the first tercile). The triple-difference coefficients θ_k identify how the treatment effect declines in each of the remaining terciles.

In a second model, I interact the DD term with terciles of the distribution of the LTV ratio in the periods prior to the Dunne judgment. High LTV ratios have three effects on default. First, in the double-trigger model, there is a pure causal effect of a high LTV ratio: borrowers default because they believe that their property is now a bad investment. Second, there is a friction to refinancing: with a leverage ratio above 100%, a new mortgage is in effect a large unsecured loan and supply is therefore rare. In that case, a high LTV ratio acts as a constraint because the borrower is unable to refinance to avoid default. This is less important in Ireland, as monetary pass-through is less

Table III
Heterogeneous Treatment Effects

This table presents the estimated coefficients from two models based on equation (4). Column (1) reports results from a model examining the liquidity effect, where the treatment effect is allowed to vary with terciles of the liquid wealth distribution. Column (2) reports results from examining the equity effect, where the treatment effect varies with terciles of the current loan-to-value (LTV) distribution. Conventional standard errors are shown in parentheses and loan-clustered standard errors are shown below them in brackets. *, ** and *** indicate significance at the 10%, 5%, and 1% level, respectively.

	Default	
	Liquidity (1)	LTV (2)
Treated × Post	0.0150	−0.0046
	(0.0053)***	(0.0052)
	[0.0088]*	[0.0051]
Post × Account Balance T2	−0.0034	
	(0.0049)	
	[0.0045]	
Post × Account Balance T3	−0.0069	
	(0.0049)	
	[0.0039]*	
Post × LTV T2		−0.0035
		(0.0050)
		[0.0046]
Post × LTV T3		−0.0063
		(0.)
		[0.0041]
Treated × Post × Account Balance T2	−0.0102	
	(0.0074)	
	[0.0103]	
Treated × Post × Account Balance T3	−0.0166	
	(0.0075)***	
	[0.0090]*	
Treated × Post × LTV T2		0.0115
		(0.0074)
		[0.0077]
Treated × Post × LTV T3		0.0198
		(0.0074)***
		[0.0087]**
Observations	7,998	7,998
Loan FE	Y	Y

complete than in the United States, for example. Finally, there is adverse selection associated with high LTV ratios at *origination*, which creates a strong correlation between the observed *current* LTV ratio and adversely selected borrowers (Gupta and Hansman (2019)). For example, in the model of Bailey et al. (2018), a borrower who wishes to maximize their exposure to housing because they have low default costs will put down a low down payment, taking a high LTV loan. This last point is an important caveat, since it is not possible to

separate the effects here. Gupta and Hansman (2019) attribute approximately 40% of the association between leverage and default to selection.

The results from this analysis are presented in Table III. Column (1) reports estimates of triple-difference regressions of the liquidity effect and column (2) reports results of regressions on the equity effect.

In column (1), the treatment effect is driven by the lowest tercile of the wealth distribution, with an estimated treatment effect of 1.5 percentage points. This result is statistically significant at the 1% level using conventional standard errors and at the 10% level using loan-clustered errors. The predicted treatment effects for the second and third terciles are both negative. The second-tercile effect is not significant, but that for the third tercile is significant and completely reverses the main treatment effect.

Column (2) examines the LTV interaction effect. The triple-difference interactions show that going from the first tercile to the third tercile of LTV (equivalent to reducing equity from an average of 40% to 111%) is associated with a 2 percentage point increase in default. This estimate is statistically significant at the 5% level when using either method to estimate standard errors.

In sum, the results show that borrower liquidity and equity are both large channels associated with the propensity to default after repossession risk is removed.

VI. Conclusion

Decreasing repossession risk reduces the cost of mortgage default, which encourages strategic defaults for both equity- and liquidity-based reasons. In this paper, I estimate the impact of repossession risk on mortgage default by examining delinquency rates in Ireland following a legal ruling that halted home foreclosures. Using DD variation from the ruling, I find that borrowers defaulted at substantially higher rates than they were otherwise likely to do. Using a sample of loans linked to borrowers' liquid account balances, I find large roles for both equity and liquidity margins.

Although the relative effect that I identify is large, the absolute effect is small. The loans studied in this paper were originated after the peak of the Irish credit bubble and are less risky on observable factors, and possibly also on unobservables. Whether 50% or 0.3 percentage points is the correct estimate for transporting the results to other scenarios depends on the functional form used to scale the treatment effects, and I do not attempt to answer this question in this paper.[20]

To my knowledge, these results are the first empirical estimates of both the moral hazard costs and the liquidity benefits of foreclosure moratoria, which were important debt relief policies in both the Great Depression and the recent Great Recession. Although I provide direct evidence on some of the important channels of ex post debt relief policies, I do not focus on the general

[20] For example, a proportional hazard model would deliver the large relative estimates, while an additive model would deliver the smaller absolute value.

equilibrium implications of the results, nor do I analyze the financial stability consequences of excess losses for lenders. If lenders expect to regularly experience elevated losses from defaulting borrowers during macroeconomic crises, then the cost of borrowing may be higher ex ante. Future research might attempt to quantify the magnitude of such equilibrium effects on future generations of mortgage borrowers.

Initial submission: January 20, 2018; Accepted: November 15, 2019
Editors: Stefan Nagel, Philip Bond, Amit Seru, and Wei Xiong

REFERENCES

Adelino, Manuel, Kristopher Gerardi, and Paul S. Willen, 2013, Why don't lenders renegotiate more home mortgages? Redefaults, self-cures and securitization, *Journal of Monetary Economics* 60, 835–853.

Agarwal, Sumit, Gene Amromin, Itzhak Ben-David, Souphala Chomsisengphet, and Douglas D. Evanoff, 2011, The role of securitization in mortgage renegotiation, *Journal of Financial Economics* 102, 559–578.

Agarwal, Sumit, Gene Amromin, Itzhak Ben-David, Souphala Chomsisengphet, Tomasz Piskorski, and Amit Seru, 2017, Policy intervention in debt renegotiation: Evidence from the Home Affordable Modification Program, *Journal of Political Economy* 125, 654–712.

Agarwal, Sumit, Souphala Chomsisengphet, Neale Mahoney, and Johannes Stroebel, 2017, Do banks pass through credit expansions to consumers who want to borrow? *Quarterly Journal of Economics* 133, 129–190.

Alston, Lee J., 1984, Farm foreclosure moratorium legislation: A lesson from the past, *American Economic Review* 74, 445–457.

Artavanis, Nikolaos T., and Ioannis Spyridopoulos, 2018, Behavioral attributes of strategic default: Evidence from the foreclosure moratorium in greece, Working paper, ssrn.

Arthur Cox, 2016, Increase in court-ordered residential repossessions in Ireland, Discussion paper.

Bailey, Michael, Eduardo Davila, Theresa Kuchler, and Johannes Stroebel, 2018, House price beliefs and mortgage leverage choice, *Review of Economic Studies* 86, 2403–2452.

Beraja, Martin, Andreas Fuster, Erik Hurst, and Joseph Vavra, 2019, Regional heterogeneity and the refinancing channel of monetary policy, *Quarterly Journal of Economics* 134, 109–183.

Bolton, Patrick, and Howard Rosenthal, 2002, Political intervention in debt contracts, *Journal of Political Economy* 110, 1103–1134.

Campbell, John Y., and Joao F. Cocco, 2015, A model of mortgage default, *Journal of Finance* 70, 1495–1554.

Campbell, John Y., Stefano Giglio, and Parag Pathak, 2011, Forced sales and house prices, *American Economic Review* 101, 2108–2131.

Citizens Information, 2016, Enforcement of debt judgments, Technical report.

Citizens Information, 2018, Credit history - the Irish Credit Bureau, Technical report.

Collins, J. Michael, and Carly Urban, 2015, When the cat's away: Payment behavior during a foreclosure moratorium, Working paper, Montana State University.

DeFusco, Anthony A., and John Mondragon, 2020, No job, no money, no refi: Frictions to refinancing in a recession, *Journal of Finance* 75, 2327–2376.

Foote, Christopher L., and Paul S. Willen, 2018, Mortgage-default research and the recent foreclosure crisis, *Annual Review of Financial Economics* 10, 59–100.

Fuster, Andreas, and Paul S. Willen, 2017, Payment size, negative equity, and mortgage default, *American Economic Journal: Economic Policy* 9, 167–191.

Gabriel Stuart, Iacoviello Matteo, and Lutz Chandler, 2020, *The Review of Financial Studies* A Crisis of Missed Opportunities? Foreclosure Costs and Mortgage Modification During the Great Recession, http://doi.org/10.1093/rfs/hhaa059.

Ganong, Peter, and Simon Jäger, 2018, A permutation test for the regression kink design, *Journal of the American Statistical Association* 113, 494–504.
Ganong, Peter, and Pascal Noel, 2020, Liquidity versus wealth in household debt obligations: Evidence from housing policy in the Great Recession, *American Economic Review* 110, 3100–3138.
Geanakoplos, John, 2010, The leverage cycle, *NBER Macroeconomics Annual* 24, 1–66.
Gerardi, Kristopher, Kyle F. Herkenhoff, Lee E. Ohanian, and Paul S. Willen, 2017, Can't pay or won't pay? Unemployment, negative equity, and strategic default, *Review of Financial Studies* 31, 1098–1131.
Ghent, Andra C., and Marianna Kudlyak, 2011, Recourse and residential mortgage default: Evidence from U.S. states, *Review of Financial Studies* 24, 3139–3186.
Goolsbee, Austan, 2014, Comments and discussion, Brookings Papers on Economic Activity Fall 2014, 119–136.
Guiso, Luigi, Paola Sapienza, and Luigi Zingales, 2013, The determinants of attitudes toward strategic default on mortgages, *Journal of Finance* 68, 1473–1515.
Gupta, Arpit, and Christopher Hansman, 2019, Selection, leverage, and default in the mortgage market, Working paper, ssrn.
Ho, Daniel E., Kosuke Imai, Gary King, and Elizabeth A. Stuart, 2011, MatchIt: Nonparametric preprocessing for parametric causal inference, *Journal of Statistical Software* 42, 1–28.
Hurst, Erik, Benjamin J. Keys, Amit Seru, and Joseph Vavra, 2016, Regional redistribution through the U.S. mortgage market, *American Economic Review* 106, 2982–3028.
Imbens, Guido W., 2004, Nonparametric estimation of average treatment effects under exogeneity: A review, *Review of Economics and Statistics* 86, 4–29.
Imbens, Guido W., and Donald B. Rubin, 2015, Causal Inference in Statistics, Social, and Biomedical Sciences (Cambridge University Press, Cambridge).
Irish Banking Federation, 2010, IBF/PWC mortgage market profile, Q2 2010, Discussion paper.
Laufer, Steven, 2018, Equity extraction and mortgage default, *Review of Economic Dynamics* 28, 1–33.
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, 2830–2857.
Melzer, Brian T., 2017, Mortgage debt overhang: Reduced investment by homeowners at risk of default, *Journal of Finance* 72, 575–612.
Mian, Atif, Amir Sufi, and Francesco Trebbi, 2015, Foreclosures, house prices, and the real economy, *Journal of Finance* 70, 2587–2634.
O'Malley, Terry, 2018, The impact of repossession risk on mortgage default, Research Technical Paper 01/RT/18, Central Bank of Ireland.
O'Neill, Pamela, 2011, Decision of Ms. Justice Dunne of 25th July, 2011 on the interaction of the Registration of Title Act, 1964, and the Land and Conveyancing Law Reform Act, 2009, Discussion paper Eversheds O'Donnell Sweeney.
Piskorski, Tomasz, and Amit Seru, 2018, Mortgage market design: Lessons from the Great Recession, *Brookings Papers on Economic Activity* 2018, 429–513.
Piskorski, Tomasz, and Vikrant Vig, 2010, Securitization and distressed loan renegotiation: Evidence from the subprime mortgage crisis, *Journal of Financial Economics* 97, 369–397.
Schelkle, Thomas, 2018, Mortgage default during the U.S. mortgage crisis, *Journal of Money, Credit and Banking* 50, 1101–1137.
Skilton, Robert H., 1943, Mortgage moratoria since 1933, *University of Pennsylvania Law Review and American Law Register* 92, 53–90.

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher's website:

Replication Code.