01/RT/18

Páipéar Taighde Teicniúil
Research Technical Paper

The Impact of Repossession Risk on Mortgage Default

Terry O'Malley

Banc Ceannais na hÉireann
Central Bank of Ireland
Eurosystem

Non-Technical Summary

This paper evaluates the claim that reducing repossession risk for homeowners leads to an increase in mortgage default. Economic theory predicts that borrowers will be more likely to default on their mortgages if their homes cannot be repossessed. After the financial crisis, commentators frequently cited the lack of repercussions as one of the contributing reasons for high mortgage arrears in Ireland.[1]

I evaluate this claim by examining how mortgage arrears evolved during the recent "Dunne judgment" period in Ireland. The legal judgment effectively removed the ability of banks to repossess homes in the event of mortgage default. The terms of the judicial decision meant that nearly every mortgaged household in the country could no longer lawfully have their homes repossessed from mid-2011 to 2013. Crucially for the analysis in this paper, a group of mortgaged households experienced no change in their repossession risk: they were exempt from the ruling of Justice Dunne. This aspect of the ruling allows to me to conduct a quasi-experimental evaluation of the impact of removing repossession risk on mortgage default: default rates for borrowers who had their repossession risk removed are compared to similar borrowers who experienced no change in the repossession regime in the event of default.

Though not the ultimate cause of the arrears crisis, I find that the removal of repossession risk led to an immediate increase in mortgage default for affected borrowers. Borrowers experiencing very low or negative levels of home equity are the most likely to default in response to the removal of repossession risk. However, this notion of purely strategic default is moderated by evidence that these "strategic defaulters" were also more likely to be in financial difficulty before the ruling. They were more likely to have missed a payment before the Dunne judgment, have lower incomes and also face higher interest rates on their mortgages.

Policy implications are straightforward. Impediments to home repossession by banks reduce a borrower's incentive to fulfil the terms of their mortgage. While a policy aiming to reduce repossession risk may benefit borrowers, it would also increase the mortgage default rate. When considering changes to repossession law, policy-makers must trade off the benefits from lower home repossessions with the moral hazard cost I have identified in this work.

[1] See for example remarks made by former Governor of the Central Bank Patrick Honohan in two 2013 speeches: https://www.centralbank.ie/news/article/remarks-prepared-by-governor-patrick-honohan-for-the-fmc2-conference-dublin ; https://www.centralbank.ie/news/article/adverse-selection-and-moral-hazard-in-forecasting-and-limiting-arrears-and-loan-losses-on-mortgages .

THE IMPACT OF REPOSSESSION RISK ON MORTGAGE DEFAULT.

Terry O'Malley[†]
Central Bank of Ireland
University College Dublin

7th February 2018

Abstract

This paper studies the effect of removing repossession risk on a borrower's decision to default on their mortgage. I exploit quasi-experimental variation in home-repossession law generated by a legal ruling in Ireland, which retroactively removed repossession risk on properties mortgaged before a certain date. Using matched data, sampled locally around this cut-off date, and a difference-in-differences research design, I find that the removal of credible repossession risk led to an immediate increase in mortgage default of 0.5 percentage points a quarter. Consistent with economic models of mortgage default, I find that the effect is driven by borrowers with low and negative home equity, but also by those with lower incomes and higher interest rates.

[†]terry.omalley@centralbank.ie. I thank Morgan Kelly and Fergal McCann for detailed comments throughout this project. I also acknowledge helpful comments from David Byrne, Paul Devereux, Andreas Fuster, Robert Kelly, Philip Lane, Michael McMahon, Rachael Meager, Conor O'Toole, Kevin Schnepel, Johannes Stroebel, Karl Whelan as well as participants at the 2017 Irish Economic Association annual conference; 2017 European Economic Association Annual Congress; 2017 UChicago/Northwestern Empirics and Methods in Economics Conference and participants at both Central Bank of Ireland and University College Dublin seminars. I also thank an anonymous reviewer for comments on this draft. Special thanks to John Joyce for assistance with the data. These views are my own, and do represent the views of the Central Bank of Ireland.

1 Introduction.

This paper evaluates the effect of repossession risk on mortgage default. The link between repossession risk and mortgage default is straightforward in economic theory. Removal or significant lowering of repossession risk amounts to a reduction in the costs of default to a borrower. Under certain conditions, borrowers may now find it optimal to default. The counterfactual scenario is that borrowers who default in the absence of credible repossession risk would otherwise not have done so. In this paper, I empirically analyse whether this theory holds true. If so, policies that inhibit home repossession in a bid to reduce foreclosures will incur a significant moral hazard cost. Owing to an unusual and unexpected legal ruling in Ireland, the repossession regime of existing mortgage holders was retroactively altered on the 25th July 2011. Crucially for this paper, there was a cut-off date specified in the ruling: properties mortgaged before 1st December 2009 could no longer be repossessed in the event of default. As this cut-off date was essentially random, borrowers who took out mortgages in late 2009 and early 2010 had their repossession regimes randomly altered one year and a half later. Using individual loan-level data, I analyse the default outcomes of two groups of borrowers separated by this cut-off date. I find that borrowers responded to the legal intervention by defaulting when they otherwise were less likely to.

My research design relies on quasi-experimental variation in which borrowers were affected by the law. Borrowers with mortgages originated before 1st December 2009 form the treatment group and those borrowers with mortgages originated after constitute the control group. I analyse matched treatment and control groups, sampled locally around this cut-off date. In the spirit of a regression discontinuity design (RDD), I limit the sample to loans originated 180 days either side of the 1st December 2009 cut-off date. I do so in order to limit potential confounding of the treatment. For example, it seems unlikely that comparing a treatment loan originated in 2006 to a control loan originated in 2010 is appropriate. The loan originated in 2006 is likely to be higher risk, due to looser credit standards in this time period. Such confounding variation would be hard to exclude even while controlling for observables, so I estimate the models locally around the cut-off. However, unlike an RDD, the variation I exploit is not cross-sectional but longitudinal: within-borrower variation before and after the judgment, between treatment and control groups. Before estimating the models, I also pre-process the data by performing matching across treatment and control groups and eliminating treatment and control observations with no similar counterpart.

Using a panel difference-in-differences estimator, I find that removing repossession risk does lead to an increase in mortgage default. I find a statistically significant average increase in default probability of 0.5 percentage points per quarter. When scaled by the control group outcome, it translates to roughly a 40% increase over the counterfactual default rate. I also identify the timing of the increase in default. Loans with a history of missed payments are the first to default in response to the judgment, while loans with no prior arrears histories respond with a slight lag.

To ensure that the research design does not either introduce or fail to eliminate bias in the treatment coefficient by construction, I also implement a randomisation inference procedure. This procedure re-tests the null hypothesis of no effect by comparing the estimates to placebo estimates obtained by permuting the cut-off date of the natural experiment. It is equivalent to a systematic and repeated robustness test of many placebo interventions, under which the null hypothesis of no effect is known to be true. The results are similar to those obtained via asymptotic inference, and indicate statistically higher default rates in the treatment group only after the time point of the Dunne judgment.

Though the average effect is informative, it is likely that only liquidity constrained borrowers or those with very low or negative levels of home equity are those who potentially benefit the most from defaulting when repossession risk is removed. Therefore the effect is likely to be driven by borrowers with these characteristics. Using new methods from the machine learning and program evaluation literature, I find that a loan's current loan-to-value ratio - the measure of a borrower's current equity position - is highly predictive of whether borrowers default in response to the removal of repossession risk. I also find that income at origination is negatively associated with the treatment effect; interest rates are positively associated with this response to the treatment. My results therefore are consistent with the theoretical model recently presented by Ganong and Noel (2016), whereby borrowers only find it beneficial to default once in very negative levels of home equity, due to utility costs of default. My results confirm that substantially reducing these default costs leads borrowers to default at lower levels of negative equity.

The rest of this paper is as follows. Section 2 discusses theories of mortgage default and introduces the natural experiment. Section 3 discusses the empirical research design and data used; section 4 introduces the regression specification; Section 5 presents the broad results and Section 6 discusses threats to identification and alternative modes of inference. Section 7 presents an analysis of heterogeneous treatment effects and external validity.

2 Repossession Risk and the Economics of Mortgage Default.

Theories of mortgage default can usually be divided into economic and social theories. Economic theories, which have received most attention (see Gerardi et al. (2017) for a recent overview), usually consider that borrowers weigh up the relative benefits and costs to paying their mortgage and default when the benefits no longer exceed the costs. For example, Campbell and Cocco (2015) present an economic model of default with fully optimizing agents. In contrast, social theory of default considers that agents perceive there to be non-economic reasons for default. In this view, mortgage default is not necessary rational. Moral and social stigma, and perceptions of fairness determine a borrower's attitude to default. Guiso et al. (2013) provide evidence that strategic mortgage default is inherently linked with borrower morality .

The two theories are not mutually incompatible. Moral aversion to default can just be incorporated into the agent's loss function, after all. But the connection between repossession risk and mortgage will depend on whether rational choice or moral considerations hold more sway on the borrower's decision to default.

According to the economic model, borrowers experience both costs and benefits to paying their mortgages each month. The benefit to paying is straightforward: the borrower can continue living in the home and amortising the loan. At the conclusion of the mortgage contract, the borrower now possesses all of the equity in the property. The cost to paying is the opportunity cost of the loan installment amount. The same money could be used to purchase goods and services, or otherwise used elsewhere. According to this theory, a borrower will default on their mortgage when the benefit of paying no longer exceeds the costs.

Economic theory incorporates the 'double-trigger' model of mortgage default (Gerardi et al., 2017). According to the double-trigger theory and its variants, a borrower will default on their mortgage under two conditions. The first is that the amount of the mortgage exceeds the value of the underlying asset. This situation is referred to as negative home equity. Due to falls in house prices, the borrower now owes the bank more money than the house underlying the mortgage is worth. Purely strategic borrowers - those who do

not pay attention to non-financial costs such as moral and social stigma- will then stop paying the mortgage. A second scenario is when the borrower can no longer afford to pay the mortgage installment due to low liquidity: the borrower does not possess the cash to actually pay. If a borrower has lost their job or had their hours reduced, then the opportunity cost of the money needed to pay the mortgage installment in full may simply be too high. A borrower might surmise that putting food on the table is more of a priority than paying back a bank. Especially when they believe the future discounted value of the home is no longer worth the sacrifice of paying a large mortgage today. This is the essence of the double-trigger theory: negative equity is necessary, but the sufficient condition for mortgage default is a binding liquidity constraint.

In a simple economic model of mortgage default, the relationship between repossession risk and mortgage default is unambiguous. Low repossession risk in the event of default reduces the costs associated with mortgage default. Absent significant repossession risk, a borrower might find it optimal to default, when in the counterfactual scenario of certain repossession risk, they would not have. The most obvious cost associated with default is the need to find alternative accommodation. Defaulting on the mortgage without the need to enter the rental market may be a powerful proposition to a borrower facing cash-flow problems or negative home equity. In this way, removing repossession risk may lower the boundary at which an agent chooses to default on their mortgage.

The social theory, in contrast, asserts that borrowers may not necessarily weigh up the economic costs and benefits to default. Moral and social stigma may make a borrower unlikely to default on their mortgage. Guiso et al. (2013) document evidence that strategic default in the US not only relates to the financial situation of the borrower but also to these moral and social considerations. Some borrowers find it immoral to default, whilst others only choose to default if they know others who have defaulted.

If the social theory is more accurate, then borrowers may not change their default decisions in response to a reduction in repossession risk. The social and moral costs still prevail over the benefit to defaulting. However, if the economic theory is more precise, then low repossession risk alters the borrower's decision to default and leads to a higher level of default.

With the exception of Guiso et al. (2013), the empirical literature has mostly focused on the economic determinants of mortgage default. Ghent and Kudlyak (2011) find that US states allowing recourse to debts have lower probabilities of mortgage default, a finding which is consistent with evidence I present in this paper. Perhaps the closest study to mine is that of Mayer et al. (2014), who find that during the great recession, US borrowers strategically defaulted in order to qualify for a mortgage modification program. Collins and Urban (2015) study the same research question as mine, but find no effect of an eight month foreclosure moratorium on default. Though they find positive effects on the rate of loan cures, this result may be mechanical due to lengthened foreclosure proceedings.

On the theoretical literature, in addition to Campbell and Cocco (2015), Schelkle (2014) presents a model where optimizing behaviour leads to double-trigger outcomes. Ganong and Noel (2016) present a model of mortgage default whereby borrowers only find it optimal to default once in very negative levels of home equity, due to costs associated with default. Lowering these costs results in an increased sensitivity of default to negative equity. My results are consistent with this model and offer empirical evidence in favour of their model.

2.1 The Dunne Judgment 2011: Legal Lacuna Removes Repossession Risk.

Ireland provides an unusual opportunity to study the effects of credible repossession risk on strategic mortgage default, through a natural experiment that removed the ability of banks to repossess mortgaged properties. The starting point of this experiment is the Land and Conveyancing Law Reform Act (henceforth Land Act), enacted in Ireland in December 2009. The 2009 legislation replaced older law from 1964 regarding the transfer of property deeds. One year and a half later, several repossession cases were being heard in the High Court. On the 25th July 2011, Ms Justice Dunne ruled that the 2009 Land Act had failed to preserve the terms of the older 1964 law that allowed lenders to repossess properties covered by mortgages (Gartland, 2011). However, a key feature of the ruling was that this legal 'lacuna' only applied to mortgages that were issued before the date of the Land Act: 1st December 2009. Mortgages originated after that date could in fact still be repossessed in the event of default. For loans issued around the time of the Land Act, the Dunne judgment quasi-randomly removed repossession risk one and a half years later. Legislation to repair the loophole in the law was not brought forward until March 2013, nearly two years later.

3 Does Removing Repossession Risk Increase Incidence of Mortgage Default?

3.1 Research Design.

The terms of the Dunne judgment in Ireland in 2011 created a natural experiment to evaluate the impact of removing repossession risk on mortgage default. The judgment was delivered in July 2011 and only applied to loans originated before the 1st December 2009, the date of the Land and Conveying Act. Loans originated before this date were no longer subject to the threat of repossession after July 2011; those originated after the cut-off date still were. The outcome of partitioning an otherwise similar set of borrowers into two such groups is naturally analysed using a difference-in-differences (DiD) estimator.

The DiD estimator estimates the causal effect of a treatment by comparing the outcomes across treatment and control groups pre- and post-intervention. The counterfactual outcome for the treatment group is estimated using the change in the outcome for the control group: what the change in outcome for the treatment group would have been, absent any intervention. Time and group fixed effects control for any unobserved confounding factors over time and within group. The key identifying assumption therefore is that the treatment effect is not confounded by any group-time varying factors. This estimator allows us to rule out confounding causes of mortgage default such as macroeconomic shocks or higher leverage in the treatment group .

The internal validity of any DiD study relies on the absence of these group-time confounding factors. The DiD model also conditions on factors within groups, so the only possible confounders are those that vary at the group-time level and directly influence loan default. An example of this would be the interest rate on a loan. It is entirely possible that bank lending practices change across the sample loan origination period; while both groups may have similar levels of interest rates, the interest rates may exhibit different time paths[1]. Though observed confounders -such as interest rate- are easily handled in the difference-in-

[1]In Ireland, this is a particular concern as there exist two types of variable-rate loans. Some loans contractually track the ECB base rate ("trackers") with a fixed margin, whilst for others the margin is allowed to float, determined by the bank

differences regression, it remains possible that unobserved factors may confound the analysis. For example, the treatment group may contain more individuals who react to macroeconomic shocks by defaulting on their loans due to other debts. Such a situation may arise when comparing a loan issued in 2007 at the height of the Irish credit boom to one issued in 2010. Even after controlling for interest rates, equity position and outstanding balance, the 2007 borrower may be more likely to default on their mortgage due to other debts built up during the boom period e.g. on a new car purchase. To combat this potential criticism, I limit the sample to loans originated in a short period around the cut-off date specified by the Dunne judgment: 180 days either side. This is not a cost-less choice. The trade-off is between estimating an unbiased effect and having a sufficiently large sample to achieve a sufficient power of statistical tests. This is similar to the regression discontinuity design, but the variation I leverage to estimate the treatment effect is longitudinal between groups created by the threshold, rather than cross-sectional.

3.2 Data.

The data are a panel data set of loans from the Central Bank of Ireland's loan-level data, a database collected from the major Irish lending institutions as part of the 2011 Prudential Capital Assessment Review. The data cover four major institutions, representing about two thirds of the Irish mortgage market. The panel runs quarterly from 1 year pre-judgment to 1 year post: eight quarters in total from October 2010 to July 2012. As part of the research design, I limit the loan-origination dates to 180 days either side of the cut-off date specified in the Dunne judgment. The data set is composed of 80,272 observations. Of 7,913 loans, 4,488 are no longer subject to repossession after the Dunne judgment. All loans are 'single-facility', meaning there are no secondary loans attached to them. The database is recorded at the loan level and contains information on loan, borrower and property characteristics including performing status, outstanding balance, current and origination loan-to-value ratio, borrower income at origination, current interest rate and rate type and origination date, among others. To reduce the chance of biasing the estimates, I use a matching algorithm to make the two groups comparable on observables at the first time point of observation.[2] I implement a matching procedure to reduce potential imbalance in the treatment and control groups in the first time period, before the treatment event. I first estimate the propensity to receive treatment as a logistic function of a range of available controls that are likely to have an impact on loan default.[3] I then match observations across the treatment and control groups using the estimated propensity score and a greedy matching method with no replacement to eliminate treatment observations with no control counterpart and vice versa (Imbens and Rubin, 2015).[4]

("standard variable"). Different compositions of these loans could lead to vastly different time varying interests rates across the two groups.

[2]The results without using the matching procedure are qualitatively similar.

[3]Bank, interest rate, interest rate type (fixed, variable), original loan-to-value ratio, income and NUTS region.

[4]I use the R package *MatchIt* to perform matching (Ho et al., 2011). Greedy matching entails matching observations without replacement. It 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. The algorithm therefore does not attempt to optimally match observations between groups, but is easier to implement and the loss over optimal matching is usually small (Imbens and Rubin, 2015).

Figure 1: Empirical default rates by group. Dashed vertical line indicates time of Dunne judgment ruling. Treatment groups are loans originated within six months prior to the cut-off date specified in the judgment. Control group are loans originated within six months after. Evidence of parallel pre-trends in loan transition rate. Default stocks diverge mechanically due to differences in transition levels, so transition is modelled.

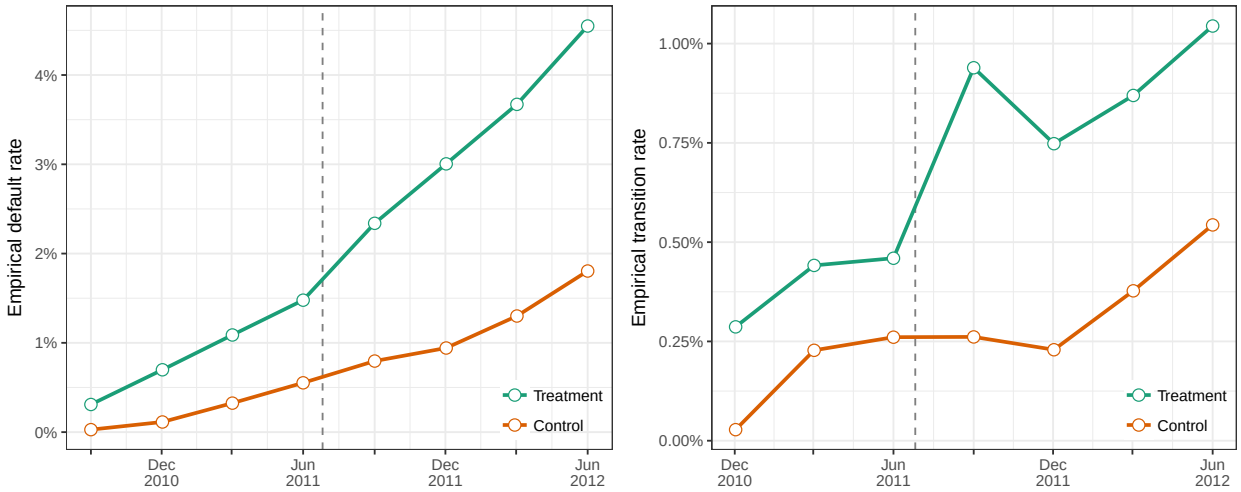

3.2.1 Graphical Analysis.

In Figure 1, I plot both the empirical default stock (level) and the associated default rate (change) for the raw data. Already it is clear that the treatment and control groups diverge after the date of the judgment: it is particularly clear from the graph of the default rate. This second panel shows the rate at which currently performing loans transition into ninety days-past-due delinquency in each quarter of the data. There is a noticeable spike in the default rate for the treatment group after the judgment. Though it subsequently falls, it remains elevated compared to its pre-judgment level. The default rate also exhibits parallel trends pre-judgment, while the default stock lines begin to diverge before the judgment. Absent any true treatment, this stock figure would begin to diverge in the case where the two groups have similar trends in their transition rates but one group has a larger level. Therefore the stock is not a good measure to use in a DiD analysis: it may be coincidental that the difference in the difference becomes statistically significant at any time point after the judgment. This statistical pattern would be entirely mechanical. The second panel provides graphical evidence in favour of a positive causal effect of the judgment.

Though there are differences in the transition rates for treatment and control groups, this does not invalidate the DiD model. As long as this difference in default levels would have remained stable in the post period, this does not undermine the causal interpretation of the results.

4 Regression Model and Statistical Inference.

To estimate the causal effect of the Dunne judgment on mortgage default, I rely on variations of the following specification:

$$\text{default}_{ibfgrt} = \alpha + c_i + \beta^{DD}(\text{Treatment}_g \times \text{Post}_t) + \mathbf{X}'_{it}\Psi + \phi_{r,t} + \tau_{b,f,t} + \epsilon_{ibfgrt} \tag{1}$$

This model is a panel data difference-in-difference estimator of the causal effect of removing repossession risk on mortgage default. I model the default probability of individual loan i; originated by bank b; with interest rate type f; in treatment group g; in region r; and in time point t. I account for differences in the composition of treatment and control groups by controlling for a matrix of loan and borrower attributes (X_{it}), as well as a set of high-dimensional fixed effects $(c_i;\phi_{r,t};\tau_{b,f,t})$ discussed below.

The dependent variable default_{ibfgrt} takes a value of 1 when a loan transitions into ninety-days past due (90DPD), the standard Basel definition of a non-performing loan. It is possible to model both thirty and sixty days past due, but these measures of default tend to be noisy: borrowers may periodically miss one or two payments due to temporary liquidity issues, or even simply due to IT issues in banking payments systems. For this reason, I use the standard Basel definition of loan default. As in section 3.2.1, I model the default rate as the *loan transition rate.*

$$\text{default}_{ibfgrt} = P(\text{90DPD}_{it} = 1 \,|\, \text{90DPD}_{is} = 0 \,\text{for all}\, s < t)$$

After default, the dependent variable is coded to drop out of the sample in $t+1$. I estimate the specification using a linear probability model, so the coefficients are directly interpreted as marginal effects. If the identifying assumptions are met, β^{DD} is the causal effect of the Dunne judgment on mortgage default. It measures the difference between the pre-judgment difference in the default rates between treated and control loans versus the post-period difference. Under the assumption that the initial difference would have carried over into the post-judgment period, β^{DD} is an unbiased estimate of the causal impact of the judgment on mortgage default.

To the question of how acceptable this identifying assumption is, I note that the treatment and control groups are matched on observable characteristics to create natural comparison groups. After matching the data, this regression specification also controls for an important set of potential time-varying confounders and a large collection of fixed effects.

c_i is a loan fixed effect, which controls for time-invariant characteristics of a loan such as the riskiness of the borrower, $\psi_{r,t}$ is a $\text{region}_r \times \text{quarter}_t$ fixed effect that controls for any regional differences in the treatment-control groups that could potentially confound the treatment effect, such as local unemployment rates. $\gamma_{b,f,t}$ is the final fixed effect and is a $\text{bank}_b \times \text{rate type}_f \times \text{quarter}_t$ fixed effect. This is a particularly powerful control, as it allows for potential group-composition differences in the bank and interest rate type (fixed v variable). If the treatment group were to feature more loans from bank A of the variable interest rate type, then changes in bank interest rate policy could be collinear with the $\text{Treatment}_g \times \text{Post}_t$ parameter. If omitted, this would be a source of bias in the estimated coefficients coming from the error term. An alternative would be to directly control for the interest rate, but this fixed effect has the advantage of being non-parametric. Therefore it is not subject to misspecification issues arising from an assumed functional form. When these time-varying fixed effects are not included then -unless otherwise indicated - I include a vector of quarter fixed effects.

X_{it} is the matrix of time-varying control variables, including the current loan-to-value ratio, interest rate (bearing in mind lack of variation when $\gamma_{b,f,t}$ is introduced into the model) and outstanding balance. In certain specifications, for example, when the loan c_i parameter is excluded, then I introduce a treatment group fixed effect and non-time varying controls such as the borrower's year of birth, income at loan origination and bank dummies.

Though the Dunne judgment offers a relatively clean natural experiment on the effect of removing repossession risk on mortgage default, there nonetheless is the possibility that the nature of the research design in this paper - using time as the treatment assignment mechanism- could lead to biases in the results, if the time-varying nature of the assignment is not accounted for. For this reason, I choose a matched treatment-control sample and estimate a flexible regression model. The estimates of β^{DD} are therefore "doubly-robust" to misspecification of either the propensity score for matching or the regression model for estimation (Robins and Ritov, 1997; Imbens, 2004). Regarding inference, I cluster standard errors at the level of individual loan, to account for within-loan correlation in the error term (Bertrand et al., 2004).

4.1 Event Study Specification.

The previous section describes the standard difference-in-differences model, in how it computes the counterfactual outcome of the treatment group, using the post-treatment outcome of the control group; and what assumptions are needed to identify the causal impact of the judgment on default. However, while this is informative, a natural analysis of the research design is to compare default rates between treatment-control groups at each time period, rather than just between the pooled pre- and post-judgment periods. The key implication of the hypothesis of this paper is that default rates should diverge only after the time of the judgment: July 2011. So, we should only observe deviations from the pre-period relationship from the September 2011 observation onward. An event study model is essentially the same as Equation 1, dissimilar in one key aspect: the $\text{treatment}_g \times \text{post}_t$ parameter is replaced by a full set of treatment and quarter interactions. This estimator allows free coefficients to be estimated for each time period's treatment effect: periods before the judgment should show no difference from the first-period difference in outcomes between treatment and control groups. Though this estimator will be less efficient than Equation 1, it provides a test of the parallel trends assumption in the standard model. Of course, parallel trends in the post period cannot be tested, but this model can provide evidence against it. I estimate variants of the following model:

$$\begin{aligned} \text{default}_{ibfgrt} &= \alpha + c_i + \sum_{t \neq Dec2010}^{Jun2012} \left\{ \beta_t \times \mathbb{1}(\text{quarter}_t) + \beta_t^{DD} \times \mathbb{1}(\text{quarter}_t) \times (\text{Treated}_g) \right\} \\ &+ \mathbf{X}_{it}'\Psi + \phi_{r,t} + \tau_{b,f,t} + \epsilon_{igt} \end{aligned} \tag{2}$$

The hypothesis of interest here is that for all quarters before July 2011, the β_t^{DD} should be statistically zero and some subset of the β_t^{DD} coefficients in the post period 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 0 by construction, and for all subsequent periods, the change in the control group's time trend is used to to estimate the counterfactual change in the treatment group. Any deviation from this prediction is considered the causal impact of the judgment. To see why this is a strict test, consider that any significant deviation from the prediction in the two periods before the judgment will show up as a non-zero treatment effect. This model therefore serves as a means of falsifying pre-trends. The identifying assumptions from section 4 carry over to here: for identification to be valid, there must be no group-time varying confounders. For this reason, this model inherits the same set of control variables and fixed effects as Equation 1. Inference is straightforward again and standard errors are clustered at the individual-loan

level. The sources of variation in this model are similar to 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 time point.

5 Estimation Results.

Estimation results of Equation 1 by ordinary least squares are presented in Table 1. Column 1 shows results from a model with only the basic ingredients of the difference-in-difference model: an intercept, a treatment group dummy, a post period dummy and the interaction of the two. The post dummy is left unreported for brevity, as are all fixed effects in subsequent models. The treatment dummy is positive and significant, meaning larger default rates for the treated loans in the periods before treatment. Treated loans are on average 0.2 percentage points (p.p.) more likely to transition to default in the pre-treatment period. This does not invalidate the identification strategy of the paper, but does suggest that the two groups are not comparable ex-ante. This simple model suggests a positive and statistically significant effect of the Dunne judgment on mortgage default. The coefficient on the Treatment $*$ Post parameter (β^{DD}) shows a 0.3 p.p. increase in default probability in the post-judgment period for the treated group. Column 2 adds time-varying controls as well as quarter fixed effects to this simple model. There are now no significant differences in the pre-period between treatment and control group default rates, once differences in composition between the two are accounted for. This is a comforting result: if both groups are comparable on outcomes in the pre-period, then it seems unlikely that they will differ in the post-period purely due to chance or an omitted variable. In this model, the treatment effect point estimate is higher at 0.4 p.p. The final column in Table 1 shows the estimates of the full model from Equation 1. Treatment is now omitted, as it is not identified when individual loan fixed effects are included. This specification also adds other high-dimensional fixed effects of $\text{region}_r \times \text{quarter}_t$ and $\text{bank}_b \times \text{rate type}_f \times \text{quarter}_t$. The treatment effect rises further to 0.5%. It is worth noting that all three treatment effect estimates are very precisely estimated, even with errors clustered at the individual loan level.

Table 1: Difference-in-differences regression. Treatment∗Post coefficient is the treatment effect of the Dunne judgment, the average increase over the estimated counterfactual in the post period. Model 1 presents a baseline model where variation is pooled across each time point pre and post. Model 2 allows for a non-parametric time trend, while model 3 controls for individual loan, bank-rate type-time and region-time fixed effects.

	Default		
	(1)	(2)	(3)
Treatment	0.002***	0.0001	
	(0.001)	(0.001)	
Treatment∗Post	0.003***	0.004***	0.005***
	(0.001)	(0.001)	(0.001)
Observations	80,667	80,663	80,663
Time-varying controls	N	Y	Y
Time FE	N	Y	-
Loan FE	N	N	Y
Region∗RateType∗Time FE	N	N	Y

*p<0.1; **p<0.05; ***p<0.01

While these estimates are informative about the causal effect of the judgment and do show evidence of strategic default in response to reduced repossession risk, the estimates of an event study model should show a positive effect of the judgment only in the post period. I therefore separately estimate the time-varying treatment effects of Equation 2 and present results in Table 2. This table shows the coefficients of the $\text{treatment}_g \times \text{quarter}_t$ parameters, divided into pre and post-judgment. Both models account for individual loan fixed effects, but model 1 does not account for time-varying controls nor the other fixed effects. The presence of individual fixed effects recovers some of the lost efficiency from separating the treatment effect into treatment-quarter effects.

The results are striking: there are no observable changes in the treatment-control group differences in either time point before the judgment. Immediately after the judgment, there is a large and statistically precise increase in the default rate for treated loans that persists for the rest of the sample time. To illustrate this treatment effect, I plot the average fitted values from model 2 for the treated group in Figure 2.[5] I also plot the estimated counterfactual outcome for the treatment group alongside: this is calculated by taking the fitted values from the model predictions, turning the $\text{treatment}_g \times \text{quarter}_t$ coefficients to 0. Intuitively, this demonstrates how the event study model is equivalent to taking the non-parametric time trend from the control group and applying it to the treatment group to obtain a counterfactual. From Figure 2, it is clear there is no difference in the pre-judgment trend between the groups, that then immediately diverges after the judgment. Given the large set of controls and high-dimensional fixed effects, along with the precise dynamics of the estimated treatment effect, it seems unlikely that there is an unobserved and powerful confounder biasing the treatment effect. This figure also contains an estimate of the increase over the counterfactual scenario. Simply computed by dividing the β_t^{DD} coefficients by the predicted counterfactual default rate, as

[5] I use the R package *lfe* to estimate the fixed effects model (Gaure, 2013). The package also allows the user to extract fitted fixed effects using the Kaczmarz method.

Table 2: Event study estimates. This specification interacts the treatment dummy with each time period for two reasons: the first is to identify the existence of pre-trends and the second is the investigate the precise timing of the change in the default transition rate. The group-differences are not statistically different from zero in both time points prior to the judgment.

	Default	
	(1)	(2)
	Pre-judgment	
Treatment∗Mar-2011	0.0005	−0.00003
	(0.001)	(0.001)
Treatment∗Jun-2011	0.002	0.0003
	(0.001)	(0.001)
	Post-judgment	
Treatment∗Sep-2011	0.007***	0.005***
	(0.002)	(0.002)
Treatment∗Dec-2011	0.007***	0.006***
	(0.002)	(0.002)
Treatment∗Mar-2012	0.008***	0.005***
	(0.002)	(0.002)
Treatment∗Jun-2012	0.008***	0.006***
	(0.002)	(0.002)
Observations	80,667	80,663
Time-varying controls	N	Y
Time FE	Y	Y
Loan FE	Y	Y
Bank∗RateType∗Time FE	N	Y
Region∗Time FE	N	Y

*p<0.1; **p<0.05; ***p<0.01

Figure 2: Estimated increase over the counterfactual for treated group. Estimates are derived from fitted values in Equation 2.

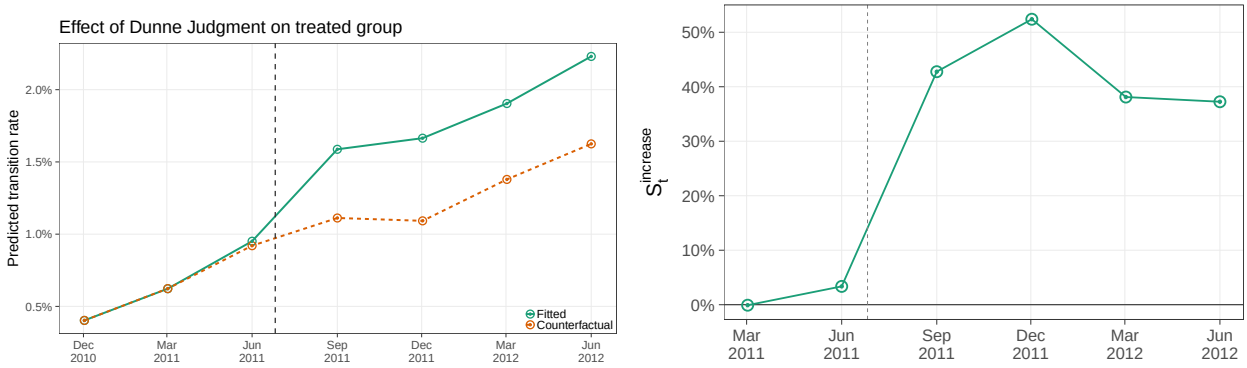

below, the results are plotted in the second panel of Figure 2 with an average value of 43% in the post-period.

$$S_t^{increase} = \beta_t^{\hat{D}D} \times \left(\frac{1}{I} \sum_{i \in \{I \cdot treatment_i\}} \left\{ \hat{c_i} + \hat{\beta}_t + \hat{\phi_{rt}} + \hat{\tau_{bft}} + X_{it}' \hat{r} \right\} \right)^{-1} \tag{3}$$

5.1 Early-Stage Arrears.

In this section, I decompose the sample into loans which were performing in the pre-period and loans with a history of one missed payment. According to my criteria, the latter group are also considered performing, as they did not enter into default. The first group must be performing in both quarters immediately prior to the judgment; they have no history of missing a payment. The second group comprises loans which have a missed payment, but did not transition to a worse arrears state between the two quarters. The reason for this second group is that I want to construct a sample of loans, where the borrower is likely to be in financial difficulty but continued to pay the mortgage in some capacity afterward. The goal is to identify borrowers with the highest opportunity cost of paying after the judgment and to analyse whether they are the first or only group to default in the post period.

The sample is composed of borrowers who were either 0 days delinquent in both March and June 2011, or borrowers who were between 1 and 30 days past due in both March and June 2011. The second group, which I refer to as the 'arrears' group, are theoretically more likely to default in response to treatment due to their higher opportunity cost of not defaulting. This distinction creates extra treatment variation within groups. There is also a mechanical explanation for why they might be the first to default. Borrowers who are current on their payments in June 2011 cannot transition to default status in September 2011, as it takes at least three months of consecutive missed payments to transition to default.

I split the sample according to these criteria and I specify a difference-in-difference-in-differences (DDD) regression model to statistically test the patterns predicted above. I am specifically interested in testing whether there is a difference in the default transition rates for the treatment and control groups within the arrears group, versus the treatment-control difference in the no arrears groups post judgment. The regression

specification is given by:

$$\begin{aligned}
\text{default}_{igat} &= \alpha + c_i + \sum_{t \neq Dec2010}^{Jun2012} \left\{ \beta_t \times \mathbb{1}(\text{quarter}_t) + \beta_t^{DD} \times \mathbb{1}(\text{quarter}_t) \times (\text{treated}_g) \right\} \\
&+ \sum_{t \neq Dec2010}^{Jun2012} \beta_t^{arrears} \times \mathbb{1}(\text{quarter}_t) \times (\text{arrears}_a) \\
&+ \sum_{t \neq Dec2010}^{Jun2012} \beta_t^{DDD} \times \mathbb{1}(\text{quarter}_t) \times (\text{treated}_g) \times (\text{arrears}_a) \\
&+ \mathbf{X}_{it}' \Psi + \epsilon_{igt} \qquad (4)
\end{aligned}$$

This specification is a DDD regression model with a fully non-parametric time trend for each treatment-arrears group. Though allowing for such a specification will reduce the power of the model, the inclusion of individual fixed effects also recovers some of this lost efficiency. I show the results from the main variables of interest in Table 3. In this model, arrears_i is an indicator variable taking a value of 1 if loan i is in early stage arrears (1-30 days past due) and 0 otherwise.

The coefficients of interest are the β_t^{DD} and β_t^{DDD} parameters , which are the time-varying treatment effects for the no arrears group (β_t^{DD}) and the time-varying difference between the treatment effect for the no arrears group and the early arrears group (β_t^{DDD}). The default rates rise in December 2011 for the treated, no-arrears group, which is the second quarter after the judgment. The effect is 0.3 p.p. For the treated, early-arrears group, the default rises by a large 0.8 p.p. in the first quarter after the judgment; consistent with this group having the greatest incentive to pay attention to and respond to the judgment. Indeed, it would not be mechanically possible for the treated, no-arrears group to enter default in this quarter. Neither group shows a statistical difference between treated and non-treated default rates in the quarter before treatment.

6 Threats to Identification and Randomisation Inference.

By the research design, the treatment assignment mechanism is at the group level i.e. assignment to treatment status is dictated by the simple rule: mortgage originated before 1st December 2009, or not. This represents the most potent threat to causal inference in this paper.

If we consider that loans originated before or after this cut-off date are likely to suffer from similar shocks, then inference on the treatment coefficient in previous section is complicated. The error term in Equation 1 is not likely to be independent across loans.[6] This would invalidate inference and overstate the precision in the β estimates. The standard remedy for such a scenario is to cluster the standard errors at the level of the regressor of interest(Bertrand et al., 2004). However, in this paper, I am investigating only two such groups. The resulting standard errors would not provide any meaningful estimate about the uncertainty of the point estimate.

Another potential solution to this problem is to abandon appeals to repeated sampling as the basis for

[6]Recall that errors are clustered within loan, so error-dependence within loans should not be a concern.

Table 3: Triple difference regression on sample of borrowers who persist in the same mortgage arrears state in the two periods before the judgment: March and June 2011. Loans are either in 0 days past due in both March and June or are in 1-30 days past due category in both months to remain in sample. This is to identify borrowers who may have received an income shock during the early sample period but did not roll straight to default.

	Default
	Pre-judgment
Treatment∗Jun-2011	0.0001
	(0.001)
Treatment∗Jun-2011∗Arrears	−0.0001
	(0.015)
	Post-judgment
Treatment∗Sep-2011	0.001
	(0.001)
Treatment∗Dec-2011	0.002**
	(0.001)
Treatment∗Mar-2012	0.002*
	(0.001)
Treatment∗Jun-2012	0.003**
	(0.001)
Treatment∗Sep-2011∗Arrears	0.081***
	(0.015)
Treatment∗Dec-2011∗Arrears	0.028*
	(0.015)
Treatment∗Mar-2012∗Arrears	0.058***
	(0.016)
Treatment∗Jun-2012∗Arrears	−0.016
	(0.016)
Observations	71,240
Time-varying controls	Y
Loan FE	Y
Time FE	Y
Bank∗Time FE	Y
Region∗Time FE	Y

*p<0.1; **p<0.05; ***p<0.01

inference, but to instead rely on the quasi-random nature of treatment assignment itself. In the difference-in-differences context of this paper, this inference resembles a type of parametric version of Fisher's randomisation inference (Imbens and Rubin, 2015). Randomisation inference under this scenario is similar to the recent work by Ganong and Jäger (2017), who provide theoretical and simulation evidence on the effectiveness of randomisation inference in the analysis of regression kink designs.

Randomisation inference essentially relies on systematic placebo testing. It is related to a robustness check used in many modern DiD studies, where the researcher attempts to contradict their methodology by subjecting it to many 'fake' or placebo interventions and re-testing the null hypothesis, when it is known to be false.

Unlike under asymptotic inference, where the theoretical uncertainty is around the random sampling of a large population, the uncertainty in the randomisation inference framework is around the random treatment assignment. Testing a hypothesis of no effect under randomisation inference is undertaken by repeatedly permuting the treatment assignment vector; re-estimating the treatment effect under these random permutation vectors; and calculating the p-value as where the true treatment effect lies in the distribution of permuted treatment effects.

It is not immediately obvious how this framework can be applied to test the null hypothesis of no effect from the Dunne judgment on mortgage default. After all, treatment here is assigned at the group-level and there is clearly no benefit to permuting treatment assignment between two groups. Instead, I rely on the following thought experiment.

Consider that the cut-off date in the Dunne judgment was a random date, that had equal chance of being 26th June 2006, as it did being the 1st December 2009. There is nothing special about the latter date: it was merely a date when new conveyancing legislation was implemented. This date bears no relation to the validity of the natural experiment in question here. It satisfies an exclusion restriction as an instrumental variable would. As I have data on all outstanding loans in Ireland during the sample period and knowledge of their origination dates, I can re-implement the preceding analysis on many different cohorts of loans. Each of these cohorts can be separated into treatment and control cohorts by a fictitious cut-off date. For a random 999 loan origination dates, I create fictitious natural experiments and store estimates of the placebo treatments. I then compare the treatment effects of the "true" natural experiment to the distribution of permuted estimates, known as the null distribution. The p-value calculated by measuring the position of the absolute value of the true estimate in the null distribution is therefore the probability that I would observe an effect size that high, when there is in fact no true effect.

To perform the randomisation inference, I first permute the cut-off date of the natural experiment 999 times to simulate alternative treatment assignments; and as in the main section, collect the loans originated 180 days either side of this cut-off date; match treatment and controls with a propensity score estimated by logistic regression and estimate a regression model.

To test the null hypothesis of no effect, I estimate exact p-values of no effect of the treatment by re-estimating the event study model (Equation 2) on these 999 data sets[7]. The results are presented in Figure 3 and show evidence of a rejection of the null hypothesis of no effect, only after the date of treatment. Figure 3. shows the time-varying null distribution of the β_t^{DD} coefficients, as well as the true estimates in red. Figure 3 also shows the associated p-values of no effect, calculated by determining what percentile of the

[7] p-values are exact because the population under study is finite.

Figure 3: Randomisation inference. Difference-in-difference coefficients for each time point relative to December 2010 difference for 999 permutations of the Dunne judgment cut-off date to create time-varying null distributions (time point after the judgment are shown in green). Actual treatment effects are shown in red dashed lines. Exact p-values of no treatment effect in each time period are shown in right panel.

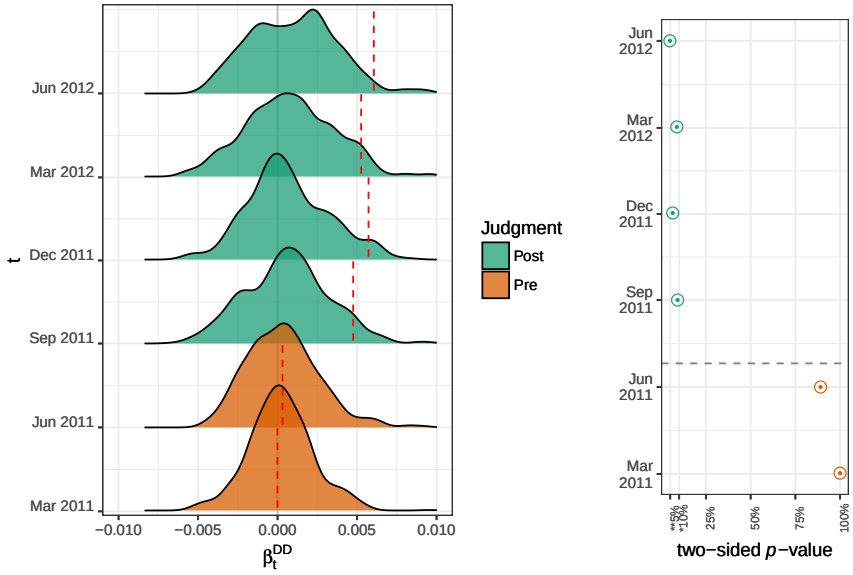

null distribution that the absolute value of the true effect lies in. For the first period (June 2010), the true treatment effect is greater than none of the placebo estimates, hence its p-value of 1. For the four periods after the judgment date, the true treatment effect lies in the range of [0.048, 0.091]. Though the p-values are higher than those estimated via asymptotic inference, they are independently significant at either the 5% or 10% level. They are therefore unlikely to have occurred by chance.

7 The Role of Leverage and Affordability in Default Response to Treatment.

So far, I have concentrated on credibly estimating the average effect of removing repossession risk on mortgage default. An analysis of how the treatment effect varies across the sample is arguably of greater interest. As discussed in Section 2, it is likely that removing such a substantial default cost changes the borrower's optimal default decision boundary. Results from Section 5 confirm that it does change, but do not reveal on what dimensions. Economic theory suggests that a borrower in negative equity or with binding liquidity constraints has the highest opportunity cost of remaining current on their payments. In this section, I attempt to empirically test this. A procedure to identify heterogeneous treatment effects across the sample is also equivalent to examining how the results might generalise to a wider population of mortgages.

To see why external validity is also important, consider that the sample is kept deliberately small so that all loans are originated close to the cut-off date. While this ensures that the sample average treatment effect is credibly identified, it also reduces the potential to generalise to the wider population of Irish mortgages, or to other mortgage markets.

A recent literature that combines machine learning algorithms with program evaluation methods has heterogeneous treatment effect estimation at its core (Athey and Imbens, 2016a,b; Foster et al., 2011; Hill, 2011; Imai et al., 2013; Hothorn and Zeileis, 2015; Seibold et al., 2016; Sivaganesan et al., 2017; Wager and Athey, 2017; Zeileis et al., 2008; Chipman et al., 2010). Since predicting missing potential outcomes is the central aim of causal inference, algorithms that optimise prediction are therefore helpful in understanding which features of a loan also predict its treatment effect. A further advantage of these methods is that they allow the analyst to estimate complex treatment response functions. Functions with many interactions and highly non-linear effects can be estimated without pre-specifying which covariates are important, or what the appropriate functional form should be.

One noted drawback of prediction algorithms however is that they are "black boxes". They are mostly used for their power to accurately predict, and not for inference about a parameter. Though Wager and Athey (2017) have recently adapted the random forest algorithm to estimate heterogeneous treatment effects, it is not so clear how to use this method to examine what factors drive heterogeneity in the treatment effect.

In this paper, I use the causal forest method to first fit a non-parametric treatment effect model, and then explore the fit of that model using simpler methods. I predict individual level treatment effects as the difference in potential outcomes for each individual, and then explore how these treatment effects vary according to loan leverage, income and interest rates using ordinary least squares regression.

The causal forest algorithm works similarly to k-nearest neighbour matching (k-nn) and kernel regression, in that treatment and control units are compared locally in covariate space. While k-nn methods break down rapidly once the number of covariates becomes large, the causal forest algorithm incorporates dimension reduction as a feature, thereby avoiding the curse of dimensionality. The causal forest algorithm is also "honest" in the language of Athey and Imbens (2016a), in that it uses a sample-splitting procedure to ensure the data used to choose the model are not those used in the estimation of the treatment effect. This method naturally produces heterogeneous treatment effects: each individual treated unit is assigned a treatment effect based on the average observed outcomes in a weighted set of local control units. Weighting is determined by the repeated recursive partitioning of the covariate space on outcomes. A single tree defines neighbourhoods in covariate space where all observations are similar on observables that predict default. Repeating this process many times allows the algorithm to establish weights: a treatment unit is compared to nearby control units if many successive trees built on different sub-samples of the data determine them to be similar on observed covariates. These weights are then used to estimate unit-level individual treatment effects. Averaging them produces an estimate of the average treatment effect. Exploring how this treatment effect function varies with covariates allows the researcher to explore what factors mattered when determining these predicted individual effects i.e. what factors are associated with treatment effect heterogeneity.

7.1 Treatment Effect Function.

To estimate heterogeneous response to treatment across the sample, I take two time points from the panel data set, $t \in \{0, 1\}$, where 0 is the quarter before the judgment and 1 is one year later (the last data point). I specify the following treatment effect model, where the dependent variable Y_i is the difference in individual i's default status from time $t = 0$ to $t = 1$:

$$Y_i = (Y_{i1} - Y_{i0}) = \tau(X_{i0},\, Z_{i1}) + \epsilon_i \tag{5}$$

The change in an individual's default status here is modelled as an unknown treatment effect function of two matrices of candidate variables and an additive error term. Note here the mixture of subscripts $i0$ and $i1$. I split variables between those measured pre-treatment and those measured after. In the former case, I am trying to minimise the "bad control" problem (Angrist and Pischke, 2008) of controlling for outcomes directly influenced by treatment.[8] In the latter case, I am trying to control for factors which may vary exogenously but are endogenously related to treatment via the research design. The big threat to inference in this case would be the interest rate: the loans originated before the treatment threshold may experience contractual changes in interest rates at a certain time point for example.

7.1.1 Causal Forest Estimation of Treatment Effect Function.

The output from the complex regression function 5 is not easily presented. The function describes individual-level treatment effects. Averaging over the sample yields an average treatment effect of 1.3 percentage points. As I am modelling one-year endline data, this is equivalent to a four-quarter effect of 0.34 percentage points, which is slightly smaller than the quarterly treatment effects estimated in Section 5, which were around 0.5 percentage points. The results may differ either due to lost efficiency in the high-dimensional method or because the non-parametric regression is able to better control for confounders.

7.2 The Role of Leverage and Affordability in Strategic Default.

The power of the estimated treatment function $\widehat{\tau(\cdot)}$ above is that it contains treatment heterogeneity by construction. How then to determine what factors, if any, are important in determining how the individual treatment effects were reached? A natural method is to compute marginal effects, as one might when presenting the results from a non-linear model such as a logit. Here, such a strategy would entail holding all covariates at their sample means and varying a parameter of interest to obtain predictions. Seeing how these predictions change with the covariate allows the researcher to understand if that parameter is important in how the treatment effects were estimated. In figure 4, I calculate the predicted treatment effect, holding all covariates except the pre-treatment loan-to-value ratio (LTV) at their means. I vary LTV between 0 and 150 and obtain predictions from the estimated treatment effect model. The model predicts a null treatment effect for low LTV ratios. This might a-priori be expected according to economic theory: the benefits from defaulting at low levels of loan leverage do not exceed the costs of default. Although noise is high in this procedure (the ribbon around the line is the estimated 95% confidence interval), the pattern at higher levels of LTV is clear: there is a sharp increase once loan-to-value begins to approach negative equity (100 LTV) and thereafter.

To understand the prediction of the model in sample, I predict the unit-level treatment effect for every observation in my data set. I then run a simple linear regression, modelling the predictions with four parameters of interest: pre-treatment loan-to-value ratio (LTV), the log of income at origination, the pre-treatment loan interest rate and the borrower's year of birth. Table 4 shows the results of the model. As expected, given Figure 4, the loan-to-value ratio is highly predictive of the individual level treatment effect. Both the log of income and the interest rate are predictive of the individual treatment effect, while the year-of-birth of the borrower has no statistical association with the predictions of the individual treatment effect.

[8] For example, the endogeneity of loan-to-value ratio and default as pointed out by Kelly and McCann (2016).

Figure 4: Predicted treatment effect as a function of pre-treatment loan-to-value ratio with 95% confidence interval ribbon. Sample average treatment effect shown in solid gray line, along with 95% confidence interval, calculated using the method of Wager and Athey (2017). Though the estimate is never statistically different from 0, the pictured relationship is illustrative of an interaction between negative equity (loan-to-value ratio $> 100\%$) and the treatment effect.

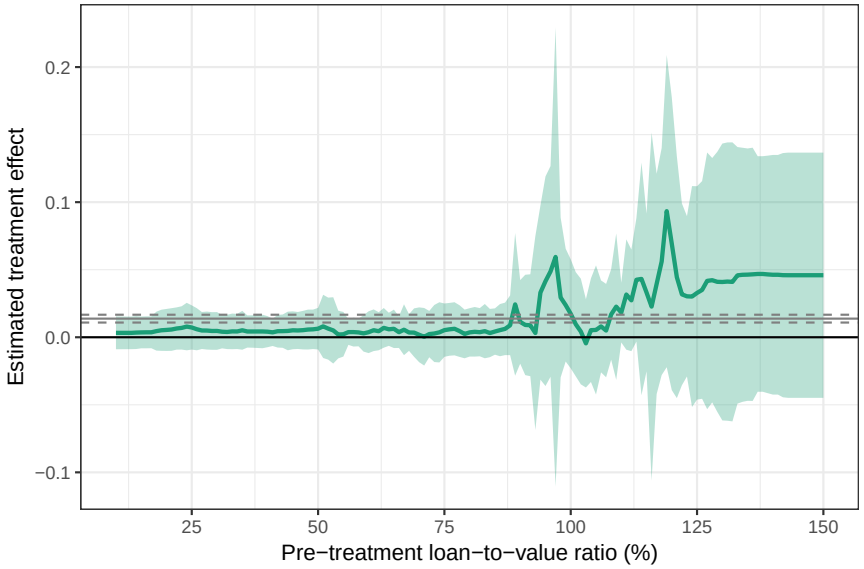

The relationship between income at origination and the treatment effect is only statistically significant at the 10% level. This relationship is likely to be noisy: current income affects borrower's current incentives and income at origination can be a noisy proxy for this.

This section provides evidence in favour of the economic model in Section 2. Not all borrowers respond to the removal of repossession risk by defaulting. Instead, only borrowers who have the most to gain from defaulting do so. Borrowers with low or negative home equity might find it optimal to default, since the future discounted value of the asset may not be worth the cost of paying the mortgage today. Borrowers with lower incomes at origination and higher interest rates may default in response to the treatment due to binding liquidity constraints and a high opportunity cost of continuing to pay their mortgage.

Table 4: Linear regression to explain predicted unit-level treatment effect.

	Treatment effect
LTV	0.0002^{***}
	(0.00001)
log(Income)	-0.002^{*}
	(0.001)
Interest rate	0.002^{**}
	(0.001)
Year of birth	-0.00001
	(0.00004)
Constant	0.027
	(0.088)
Observations	12,298

*p<0.1; **p<0.05; ***p<0.01

8 Conclusion.

This paper examines the impact of removing repossession risk on mortgage default. By analysing a natural experiment arising from an unusual legal ruling in Ireland, I have identified a moral hazard cost to policies that effectively reduce the cost of defaulting to the borrower. Borrowers respond to the removal of repossession risk by defaulting, on average, more than they otherwise would have. This effect occurs immediately and persistently. Though tempting to label as strategic default, this average effect is not only driven by the equity position of borrowers, but also by affordability concerns: measured here by the borrower's income at origination, loan interest rate and prior missed payments.

Policy implications are straightforward. Impediments to home repossession by banks reduce a borrower's incentive to fulfill the terms of their mortgage. While a policy aiming to reduce repossession risk may benefit borrowers, it would also increase the mortgage default rate. When considering changes to repossession law, policy-makers must trade off the benefits from lower home repossessions with the moral hazard cost I have identified in this work.

References

Angrist, J. D. and Pischke, J.-S. (2008). *Mostly harmless econometrics: An empiricist's companion*. Princeton university press. 7.1

Athey, S. and Imbens, G. (2016a). Recursive partitioning for heterogeneous causal effects. *Proceedings of the National Academy of Sciences*, 113(27):7353–7360. 7

Athey, S. and Imbens, G. (2016b). The state of applied econometrics-causality and policy evaluation. *arXiv preprint arXiv:1607.00699*. 7

Bertrand, M., Duflo, E., and Mullainathan, S. (2004). How much should we trust differences-in-differences estimates? *The Quarterly journal of economics*, 119(1):249–275. 4, 6

Campbell, J. Y. and Cocco, J. F. (2015). A model of mortgage default. *The Journal of Finance*, 70(4):1495–1554. 2

Chipman, H. A., George, E. I., McCulloch, R. E., et al. (2010). Bart: Bayesian additive regression trees. *The Annals of Applied Statistics*, 4(1):266–298. 7

Collins, J. M. and Urban, C. (2015). When the cat's away: Payment behavior during a foreclosure moratorium. Technical report, Working Paper. 2

Foster, J. C., Taylor, J. M., and Ruberg, S. J. (2011). Subgroup identification from randomized clinical trial data. *Statistics in medicine*, 30(24):2867–2880. 7

Ganong, P. and Jäger, S. (2017). A permutation test for the regression kink design. *Journal of the American Statistical Association*, (just-accepted). 6

Ganong, P. and Noel, P. (2016). The effect of debt on default and consumption: Evidence from housing policy in the great recession. 1, 2

Gartland, F. (2011). Mortgage cases face delays over law loophole. [Online; posted 10-October-2011]. 2.1

Gaure, S. (2013). lfe: Linear group fixed effects. *The R Journal*, 5(2):104–117. 5

Gerardi, K., Herkenhoff, K. F., Ohanian, L. E., and Willen, P. S. (2017). Can't pay or won't pay? unemployment, negative equity, and strategic default. *The Review of Financial Studies*, page hhx115. 2

Ghent, A. C. and Kudlyak, M. (2011). Recourse and residential mortgage default: evidence from us states. *Review of Financial Studies*, page hhr055. 2

Guiso, L., Sapienza, P., and Zingales, L. (2013). The determinants of attitudes toward strategic default on mortgages. *The Journal of Finance*, 68(4):1473–1515. 2

Hill, J. L. (2011). Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics*, 20(1):217–240. 7

Ho, D. E., Imai, K., King, G., and Stuart, E. A. (2011). MatchIt: Nonparametric preprocessing for parametric causal inference. *Journal of Statistical Software*, 42(8):1–28. 4

Hothorn, T. and Zeileis, A. (2015). partykit: A modular toolkit for recursive partytioning in R. *Journal of Machine Learning Research*, 16:3905–3909. 7

Imai, K., Ratkovic, M., et al. (2013). Estimating treatment effect heterogeneity in randomized program evaluation. *The Annals of Applied Statistics*, 7(1):443–470. 7

Imbens, G. W. (2004). Nonparametric estimation of average treatment effects under exogeneity: A review. *The review of Economics and Statistics*, 86(1):4–29. 4

Imbens, G. W. and Rubin, D. B. (2015). *Causal inference in statistics, social, and biomedical sciences*. Cambridge University Press. 3.2, 4, 6

Kelly, R. and McCann, F. (2016). Housing equity and mortgage default: Introducing the missed payment bias. *Mimeo*. 8

Mayer, C., Morrison, E., Piskorski, T., and Gupta, A. (2014). Mortgage modification and strategic behavior: Evidence from a legal settlement with countrywide. *American Economic Review*, 104(9):2830–57. 2

Robins, J. M. and Ritov, Y. (1997). Toward a curse of dimensionality appropriate(coda) asymptotic theory for semi-parametric models. *Statistics in medicine*, 16(3):285–319. 4

Schelkle, T. (2014). Mortgage Default during the U.S. Mortgage Crisis. (72). 2

Seibold, H., Zeileis, A., and Hothorn, T. (2016). Model-based recursive partitioning for subgroup analyses. *The international journal of biostatistics*, 12(1):45–63. 7

Sivaganesan, S., Müller, P., and Huang, B. (2017). Subgroup finding via bayesian additive regression trees. *Statistics in Medicine*, 36(15):2391–2403. 7

Wager, S. and Athey, S. (2017). Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, (just-accepted). 7, 4

Zeileis, A., Hothorn, T., and Hornik, K. (2008). Model-based recursive partitioning. *Journal of Computational and Graphical Statistics*, 17(2):492–514. 7