UNIVERSITÀ DELLA SVIZZERA ITALIANA

DOCTORAL DISSERTATION

Essays in Household Finance and Monetary Policy Transmission

Author:
Virginia GIANINAZZI

Supervisor:
Prof. Alberto PLAZZI

*A dissertation submitted in fulfillment of the requirements
for the degree of
PhD in Economics
SFI PhD in Finance*

in the

Institute of Finance
Faculty of Economics

Dissertation Committee:

Prof. Francesco FRANZONI, Università della Svizzera italiana

Prof. Loriana PELIZZON, Goethe University

Prof. Alberto PLAZZI, Università della Svizzera italiana

Prof. Marti SUBRAHMANYAM, NYU Stern School of Business

June 2, 2020

Abstract

The chapters in my dissertation use novel empirical settings to contribute new insights to two fundamental questions in finance: how households make financial decisions and how market prices react to uninformative demand shocks.

In the first chapter, *Reference Points in Refinancing Decisions*, I exploit the unique design of mortgages in the UK to study how households make mortgage refinancing decisions. Several recent papers show that many borrowers miss out on substantial savings by failing to refinance their mortgage when interest rates decline. Yet, we know little about why households are often inactive in response to interest rate incentives. In this paper, I identify reference dependence as an important source of inactivity. Consistent with borrowers making decisions according to prospect theory, I find that refinancing choices are significantly affected by interest rates that individuals were charged in the past. While past rates are by design unrelated to the opportunity cost of inaction, the evidence suggest that they serve as salient reference points against which borrowers define gains and losses. The effect is estimated around pre-determined dates, when mortgages automatically reset to a punitive rate unless borrowers take action and refinance to a new product at current market rates. The exogeneous timing of the refinancing decision and the absence of borrower-specific pricing of mortgages in the UK allow to identify the causal effect of reference points. The evidence that households leave substantial money on the table unless faced with out-of-pocket losses suggests that savings forgone by sticking to an expensive fixed-rate mortgage are not perceived as an actual loss, helping to explain the widely observed inertia despite falling interest rates and shedding light on a behavioral friction to the pass-through of expansionary monetary policy to households balance sheets.

In the second chapter, *Mortgage Default and Positive Equity: Lessons from Europe*, coauthored with Loriana Pelizzon and Alberto Plazzi, we study the timing of mortgage default when borrowers face the threat of recourse by lenders. The common view on lender recourse is that it reduces delinquencies and foreclosures by encouraging more responsible borrowing ex-ante and discouraging strategic default ex-post. Default is also usually described as the exercise of a real option, but this option value is nullified in the presence of fully enforceable recourse. Under lender recourse, default should not depend on the level of equity and should result only from liquidity shocks. Still, borrowers should default only if equity is negative, otherwise they would be better off by selling their house and pre-paying the mortgage. We posit that borrowers are not indifferent between liquid income and illiquid wealth in the form of housing equity, and therefore may prefer to forego their equity in the house in order to avoid facing income garnishment in case of default. We show that the majority of defaults happen when collateral would be in principle enough to repay the debt. We also find that equity at default is significantly negatively related to household income at origination, which is consistent with the threat of recourse being greater for borrowers with a higher marginal utility of consumption.

In the third chapter, *Quantitative Easing and Equity Prices: Evidence from the ETF Program of the Bank of Japan*, co-authored with Andrea Barbon, we are interested in the workings of quantitative easing and, more broadly, in the asset pricing implications of exogeneous changes in the available quantity of assets. We study the effect on the price of the underlying stocks of the large-scale purchases of equity ETFs that the Bank of Japan (BOJ) has been carrying out as part of its QE program with the intention of reducing risk premia. We run an event study around two dates when the BOJ announced an expansion of the purchase target and we find a positive, sizeable and persistent impact on stock prices. We exploit the heterogeneity of the induced shock to supply to show that the variation in event returns in the cross-section is consistent with the change in the marginal contribution of each stock to the risk of the aggregate portfolio held by private investors. This evidence is consistent with a model where QE reduces the quantity of assets held by the private sector, effectively changing the risk composition of the aggregate portfolio of the representative agent. For this to be an equilibrium, prices need to adjust to ensure market clearing, implying downward sloping demand curves. The estimated net effect of the policy is a 20 basis points increase in aggregate market valuation per trillion Yen invested into the program, corresponding to a price elasticity of 1.

Acknowledgements

A PhD is as much an individual as a team accomplishment. I would never have been able to survive the ups and downs of this journey without the support of the people who have always been at my side and of those whom I met along the way.

I am especially grateful to my dissertation advisor, Alberto Plazzi, for his guidance and encouragement. I thank him for believing in me as a researcher, for involving me in exciting research projects and for his constant support on the academic job market. Thank you, Alberto, for always making me feel like I had somebody in my corner.

I am indebted to Marti Subrahmanyam for the fantastic opportunity of spending one year at NYU. I greatly admire Marti for his dedication to the profession and his generosity towards students. The discussions with him, the seminars and the interaction with the faculty at Stern have been a great inspiration and motivation. New York city stole my heart and I am forever thankful to everyone that made this possible.

I would like to thank the remaining members of my dissertation committeee, Loriana Pelizzon and Francesco Franzoni. I am grateful to Loriana for taking me under her wing and for her positive spirit; and to Francesco whose intellectual curiosity and challenging discussions have always spurred me to work harder.

I owe a big thank you to Laurent Frésard for his energy, time and perspective. His comments helped to substantially improve my work.

This thesis also benefited from many comments by my fellow doctoral students, who also spared no effort in proofreading parts of it. I thank Andrea Barbon, with whom I shared the thrill of the first publication. I thank Roberto, Silvia, Julia and Iris for being such fantastic colleagues and friends. I thank Davide for his enthusiasm and friendship.

My dissertation is dedicated to Andrea, who makes my life beautiful. It is also dedicated to my best friend, Giulia, my parents, Renzo and Paola, my sister, Cornelia, my brother, Vanni, my grandmother Lidia, and my splendid nieces, Olivia, Alice and Bianca.

Contents

1 Reference Points in Refinancing Decisions

1.1 Introduction

In countries where the majority of mortgages are fixed rate, the ability of the central bank to stimulate consumption through the refinancing channel of monetary policy crucially relies on households optimally responding to financial incentives. However, several studies document that borrowers wait too long to refinance their mortgages when interest rates fall, thus missing out on substantial savings and imposing a friction to the pass-through of low rates onto households balance sheets (Agarwal, Rosen, and Yao, 2016; Andersen et al., 2015; Bajo and Barbi, 2018; Campbell, 2006; Keys, Pope, and Pope, 2016). Despite its key role in monetary policy transmission and its implications for household welfare, our understanding of households' refinancing decision-making is still limited. While showing that borrowers make refinancing mistakes is already a complicated task, since we generally do not observe neither the rates they were offered (if any) nor the upfront fees, showing *why* they make them is even more challenging.

One potential explanation for the observed sluggishness in mortgage refinancing is that people treat opportunity costs differently than "out-of-pocket" costs (Johnson, Meier, and Toubia, 2019; Kahneman, Knetsch, and Thaler, 1991). Foregone savings implied by expensive mortgage payments may not be perceived as an actual loss, if only deviations from regular payments are considered as a loss or a gain. Empirically, testing the hypothesis of reference dependence in refinancing choices is hard, either because individual reference points are not observable, or because actual payments never deviate from reference payments unless borrowers do refinance. To address this challenge, I exploit the design of mortgages in the United Kingdom where the interest rate on a typical mortgage is scheduled to reset after an initial period, at which point borrowers are faced with the choice of refinancing to current market rates or letting the rate automatically change to a punitive reversion rate.[1] Reversion

[1]Cloyne et al., 2019 exploit the same setting to investigate the relationship between house prices and borrowing. Interest rate resets of adjustable-rate mortgages in the US have been used as quasi experimental variation by various authors to explore different questions. Fuster and Willen, 2017 and Tracy and Wright, 2016 use data on hybrid ARMs in the US to study the impact of refinancing on default and find that reductions in mortgage payments lead to a substantial decrease in default probabilities. Di Maggio et al., 2017 similarly rely on the variation in the reset timing as an exogenous shock to income to study the real effects of decreasing debt servicing costs, but extend the analysis to consumption responses and debt overpayment.

rates are punitive because they are usually substantially higher than current rates, and increasingly so in recent years. At the same time, *expensive* reversion rates can be *cheaper* than the expiring initial rate, depending on when the loan was originated and the evolution of interest rates since then. This setting allows me to disentangle borrowers' responses to potential savings from responses to changes in mortgage payments relative to the past. I find that the probability to refinance decreases the larger the nominal gain (or the smaller the nominal loss) experienced in case of inaction. In particular, borrowers who perceive inaction as *relatively* cheap compared to the past are significantly more likely to stick with the *absolutely* expensive rate that applies by default after the interest rate resets. This negative relationship is apparent in Figure 1.1, which plots the average refinancing probability conditional on the relative gain in case of inaction. Refinancing is significantly less frequent when gains are positive. This evidence is at odds with models of optimal refinancing (Agarwal, Driscoll, and Laibson, 2013). Since past rates are uninformative about borrowing costs going forward, they should not matter for refinancing choices.

The idea that the utility of an outcome is a function of the outcome's distance from a reference point is a fundamental tenet in prospect theory (Tversky and Kahneman, 1981). Extensive evidence from laboratory experiments and field studies documents the importance of framing effects on decision making.[2] While this paper is the first to look at the effect of reference points on mortgage refinancing decisions, in finance reference-dependence has been studied in a number of other settings (Baker, Pan, and Wurgler, 2012; Loughran and Ritter, 2002), including borrowing markets and housing decisions. Closest to this paper are the results in Dougal et al., 2015, which shows from the syndicated loan market that firms borrowing rates seem unduly influenced by previous rates, providing evidence that uninformative historical information may enter negotiations through the effect of reference points. Andersen et al., 2019 and Genesove and Mayer, 2001 study anchoring and reference dependence in listing premia in the housing market and find that listed prices increase sharply when households face nominal losses. While previous work establishes the role of reference points or anchors to simplify the complex tasks of valuation and negotiation, in my setting there is no bargaining between borrowers and lenders or buyers and sellers, and there is no concern about the endogenous timing of the choice with respect to the nominal loss. Moreover, the large effects of reference points that I estimate are particularly puzzling given the large amount of money left on the table: between 2013 and 2017, households that do not refinance pay on average an interest rate 2.26 percentage points higher than current market rates.

Reference points and deviations from them arise from the design of the typical mortgage in the UK. Unlike in the United States, long-term fixed rate mortgages are not available in the UK. In fact, most borrowers are on an initial fixed rate mortgage for a short period of time (typically 2 or 5 years), after which the mortgage reverts to a variable rate for the remainder of the term. Once the fixed period ends, borrowers

[2] For a comprehensive survey of the literature see e.g. Barberis, 2013; Beshears et al., 2018; DellaVigna, 2009; Hirshleifer, 2015.

Figure 1.1: Conditional Refinancing Probability

The figure shows average refinancing probability conditional on the relative gain in case of inaction, expressed in percentage of the outstanding loan balance. This is the difference between the rate on the expired initial deal and the reversion rate (SVR) that applies after the reset. In the positive region, borrowers that do not refinance experience a decrease in monthly payments if they revert to the SVR. The relative gain is instead negative for borrowers who would experience an increase in monthly payments by reverting on the SVR at the end of the initial fixed period. For this figure, I restrict my sample to mortgages that have a LTV between 60% and 75%, and an outstanding balance between £100,000 and £250,000 at reset. This corresponds to 7,594 observations. The red bars show 95% confidence intervals for the conditional mean.

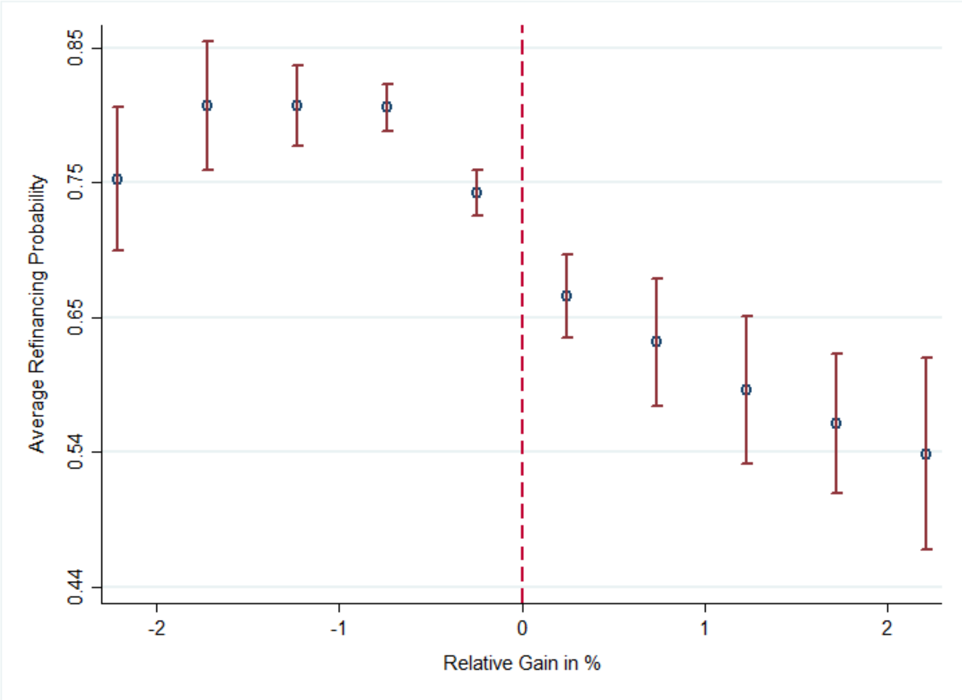

can refinance to a new initial deal without incurring a prepayment penalty. Paired with the fact that reversion rates tend to be much higher than current market rates, this creates a strong incentive to refinance around the reset date. While reversion rates and current market rates determine the savings foregone by households that fail to refinance, the *change* in mortgage payments experienced in case of inaction depends on the just expired initial fixed rate. I posit that borrowers evaluate the benefits from refinancing relative to this expired fixed rate, which is a natural candidate for a reference point in this context. I test this hypothesis using loan-level data where I can follow borrowers' refinancing behavior after the initial fixed period ends, and which allows me to observe both the matured fixed rate as well as the reversion rate that applies by default. I find that the difference between the matured fixed rate and the reversion rate is a significant predictor of the heterogeneity in refinancing decisions among borrowers, after controlling for differences in potential savings, mortgage characteristics and observable demographics.

The empirical strategy I propose in this paper leverages different institutional features of the UK mortgage market to identify the causal effect of reference points on refinancing decisions. The main advantage of using UK data is that the design of mortgages implies that the *status quo* level of payments is not preserved in case of inaction. This induces cross-sectional variation in gains and losses with respect to the reference point that is essential to test for reference dependence in how borrowers refinance. A second key feature is that in the UK there is no ex-post price discrimination based on borrower-specific characteristics, including credit scores and income. This rules out obvious endogeneity concerns about past mortgage rates with respect to refinancing opportunities. One would otherwise worry that borrowers who were paying higher rates in the past will naturally face higher rates upon refinancing as well. Crucially, mortgage rates in the UK are quoted by lenders as a discrete schedule at maximum Loan-to-Value (LTV) ratio in steps of 5 to 10% (Best et al., 2018) and apply to all eligible borrowers at a given point in time. I will therefore argue that lender $\times$ reset time $\times$ LTV-buckets fixed-effects absorb all hetereogeneity in rational refinancing incentives across borrowers, after controlling for mortgage size and remaining time to maturity. The reason for this is that borrowers who reset at the same time, with the same lender and have a similar LTV face the same reversion rate and the same set of refinancing rates.

Unlike other studies, I can rely on cross-sectional identification because there is a well-specified time window in which refinancing becomes a salient choice and where one can therefore easily compare decision outcomes across people. This would clearly not be possible in the case of pre-payable long-term fixed rate mortgages, where borrowers can refinance at any time. In such cases, researchers have to compare actual behavior with a model implied optimal benchmark, which can be hard to compute and has to rely on a number of assumptions. Instead, the cross-sectional approach in this paper is based on a simple argument. Since refinancing is optimal if and only if the saving is larger than the sum of the upfront cost and the option value of refinancing in the future, borroweres with comparable mortgage debt, who face the same rate in case of inaction, the same available market rates and the same fees, should all optimally either refinance or stick with the reversion rate. Showing that there are systematic differences in refinancing behavior predicted by backward-looking information provides evidence of reference dependent decisions. Importantly, this statement holds regardless of whether refinancing is actually optimal or not.

In my analysis, I assume that the rate on the initial fixed period is the relevant reference point for refinancing decisions. Anecdotal evidence suggests that people do take their current fixed rate into consideration when they evaluate the benefits from refinancing. Numerous articles in the popular press warn borrowers about the possible jump in mortgage payments at the end of the fixed period.[3] An article in the Financial Times, 2017 even refers directly to the large difference between the reversion rate and the maturing fixed rate as a determinant of refinancing incentives ("So there's motive for people to remortgage? Precisely."). However, while the status quo

[3]"Every month hundreds of thousands of borrowers reach the end of their fixed-rate mortgage deal. In most cases, that means their mortgage payments are set to rise - in some cases by a lot. But you can take action to avert these higher costs." (The Telegraph, 2019)

seems to predict people choices in many settings, including this one, it is less clear what should determine reference points in theory. Kőszegi and Rabin, 2006; Kőszegi and Rabin, 2007 argue that expectations determine reference points and that the status quo only matters when people expect to preserve it in the future. In the context of this paper, it is hard to tell whether the expired rate matters through the current level of payments or because people extrapolate current borrowing costs into the future. Since the reset of the mortgage rate happens on a pre-determined date and reversion rates are observable over time, the change in mortgage payments is predictable. However, it is still plausible to think that borrowers did not budget for the predictable change in payments and would find themselves forced to cut consumption unless they manage to refinance their mortgage. An alternative explanation is that the sudden jump in payments serves as a wake-up call for borrowers, who will check the current level of interest rates if and only if the interest rate increase. Otherwise, if interest rates decrease, borrowers will not make the effort of looking at the new available rates and will not realize that savings can be made.

The analysis faces three main identification challenges due to the lack of random assignment of reference points. The variation in past rates that I use to estimate the effect of reference dependence on refinancing decisions comes primarily from differences in borrowers' choices about the length of the initial fixed interest rate period. While this would not be problematic in general, during the sample period there is a strong positive correlation between past rates and the length of the fixation period because of steadily falling mortgage rates since the financial crisis. The first concern is that unobserved borrower characteristics that simultaneouly explain both a preference for less duration risk and a lower propensity to respond to financial incentives might be driving my results (Koijen, Van Hemert, and Van Nieuwerburgh, 2009). I present a battery of results to show that confounding unobserved heterogeneity in preferences for initial deal duration are unlikely to explain my findings. First, I show that at least in terms of observable characteristics, borrowers that at the same point in time choose different fixation periods are largely similar. Then, I use data from the BOE/NMG Survey of Household Finances to shows that in a period where longer fixation did not imply higher reference rates, borrowers that are expected to have a preference for bearing less duration risk were refinancing more frequently than the average borrower. Lastly, from a placebo regression on a subsample of loans that reset with similar past rates, I demonstrate that initial mortgage duration has no significant effect on refinancing probabilities. The second concern is that borrowers that choose shorter fixation periods are faced with refinancing decisions more frequently. This might both introduce survivorship bias and make borrowers on shorter fixation periods more experienced due to learning through repeated refinancing. I address both issues looking at borrowers that are faced with a rate reset for the first time. I show that the effect of reference points on refinancing decisions is still strong in a subsample where survivorship is ruled out and borrowers are expected to be equally experienced. Third, an alternative explanation for my findings is that borrowers that select into longer maturities face a higher probability of being denied refinancing. This is a concern given the results in Hertzberg, Liberman, and Paravisini, 2018, which shows from the peer-to-peer lending market in the US that borrowers with higher unobservable repayment risk tend to self-select into longer maturity contract based on private information. I show that results do not change when I control for the incentive to self-select based on unobservables using the estimated difference in term-premia between 2 and 5-years maturities at the time of the origination of the mortgage.

Refinancing decisions play a key role for the effectiveness of monetary policy in stimulating aggregate consumption by reducing the cost of debt servicing. Reflecting this policy importance, there has been a surge of papers in recent years that investigate frictions to refinancing. After accounting for the effect of negative equity (Beraja et al., 2018; Agarwal et al., 2015b), upfront costs and documentation requirements (DeFusco and Mondragon, 2018) in inhibiting refinancing especially during recessions, a number of papers document that households do not refinance optimally (Andersen et al., 2015; Agarwal, Rosen, and Yao, 2016; Bajo and Barbi, 2018; Campbell, 2006; Johnson, Meier, and Toubia, 2019). In particular, Keys, Pope, and Pope (2016) show that more than 20% of households in the US are paying too much for their mortgage, incurring a median loss of more than $10,000 in present value terms. Some papers then investigate the determinants and the heterogeneity of sluggishness in refinancing. Using Italian data, Bajo and Barbi (2018) find that this "financial apathy" is strongly related to socio-demographic characteristics and household financial literacy. Andersen et al. (2015) use loan-level data on mortgages in Denmark to try to quantify the relative importance of two sources of inactivity, namely inattention and inertia. While inertia is supposed to disappear when interest rate incentives are sufficiently large, inattention can prevent people from refinancing even when the incentive to do so is strong. The paper exploits the difference in implied refinancing dynamics to quantify the relative importance of these two channels. While both drivers appear to be important, inattention seems to be the main determinant of low refinancing among households with a low socio-economic status. Johnson, Meier, and Toubia (2019) analyze administrative data on pre-approved offers and argue that time preferences and lack of trust are leading factors explaining the low refinancing rates observed in their sample. While suspicion towards financial institutions seems to be one of the motives that prevent households from refinancing, Maturana and Nickerson (2018) show that peer effects can strongly increase refinancing rates. The results in my paper make several contributions to this literature. First, my paper is the first to provide empirical evidence that the fact that missing out on a saving opportunity does not constitute a nominal loss significantly decreases households' propensity to refinance. Second, my results imply that the responsiveness of a borrower not only depends on individual attributes such as financial literacy, but is crucially affected by the "framing" of the refinancing gains. Third, a lot of the engagement in the refinancing market in the UK seems to be motivated by the desire of avoiding a nominal loss, and this can be easily misinterpreted as a sign of financial sophistication. This could lead to wrong estimates about the actual responsiveness of mortgages when interest rates go down in a recession and thus overestimate the stimulating potential of expansionary monetary interventions. This paper fits therefore more broadly in the rapidly growing literature on the role of mortgage markets and security design in the transmission of monetary policy through the refinancing channel (Abel and Fuster, 2018; Auclert, 2019; Berger et al., 2018; Di Maggio, Kermani, and Palmer, 2016; Eichenbaum, Rebelo, and Wong, 2018; Fuster and Vickery, 2014; Greenwald, 2018; Wong, 2019). I establish an important complementarity between monetary policy and the decisions of mortgage lenders in the UK about where to set reversion rates, which appear to be crucual in amplifying refinancing frictions coming from behavioral biases. In a similar spirit to Berger et al., 2018, who argue that the average outstanding rate on fixed-rate mortgages leads to a path-dependent effectiveness of monetary policy through the incentives to prepay, my results show that in the UK the effectiveness of the refinancing channel of monetary policy depends on the distribution of reference points and, more precisely, of the expected nominal gain in case of inaction. Finally, the evidence provided in this paper about borrowers' reluctance to taking action also relates to the literature on the effects of default options on economic outcomes (Beshears et al., 2009; Beshears et al., 2015) which finds, in a number of different settings, a strong tendency of people against opting out that is hard to reconcile with any plausible value of transaction costs.

The rest of the paper is organized as follows. Section 1.2 describes the institutional backround of my analysis. Section 1.3 introduces the theoretical framework. Section 1.4 presents the empirical strategy. Section 1.5 describes the data and sample I use. Section 1.6 presents the main empirical results and Section 1.7 addresses a number of identification challenges. Section 1.8 concludes.

1.2 Mortgage Design and Reference Points in the UK

Unlike in the United States, where the most common product is a 30-years fixed rate mortgage, homeowners in the UK can lock in their mortgage rate only for limited periods of time. The typical mortgage charges an initial fixed rate for a period of 2-5 years, at the end of which the mortgage automatically reverts onto the current *Standard Variable Rate (SVR)* of the lender. At the end of the initial fixation period, the borrower has the option to refinance to a new initial deal at current market rates without penalty. Borrowers rarely prepay before the end of the introductory deal since most contracts feature large early repayment fees, typically 5 percent of the outstanding loan amount (Best et al., 2018, Cloyne et al., forthcoming). At the end of the fixed rate period, the incentive to refinance is strong for most borrowers. In fact, SVRs charged by lenders are usually susbtantially higher than new fixed or variable market rates quoted at the same point in time. Reversion rates are therefore *expensive* relative to market rates and households who do not refinance might be missing out on considerable savings. Moreover, even though the SVR is a variable rate and is therefore expected to go up when interest rates increase, there is no guarantee that it will fall if interest rates decrease. This is because each mortgage lender sets its own SVR and can revise it at any time, with no obligation to follow the BOE's base rate or any wholesale rate.

This mortgage design implies that, on pre-determined dates, borrowers come off the fixed rate that they have been paying for the previous 2 to 5 years and are faced with the choice whether to refinance to a new fixed rate or to stay on the SVR of their lender. While lenders usually set their SVR above market rates, whether the SVR is above or below a given borrower's matured fixed rate at the end of her initial period also depends on the path of interest rates between the origination and the expiration of the fixed deal. This means that by reverting onto the SVR some borrowers might see their monthly mortgage payments go down. For these borrowers, the expensive SVR is therefore *cheap* relatively to their own past rate. On the contrary, for borrowers who expect their mortgage payments to go up at the end of the initial period, the SVR is expensive both relative to market rates and relative to the expired fixed rate. In the data, the distribution of matured fixed rates around the SVR varies over time and, at any given point in time, we can observe substantial cross-sectional heterogeneity.

Figure 1.2 plots average quoted 2-years and 5-years fixed rates (solid lines), as well as the average SVR applied by mortgage lenders (dashed line) from 2000 to 2017. I focus on these two maturities (fixation periods) because they are by large the most common in the UK. The cut of the base rate by the BOE at the end of 2008 led to a visible structural break in the relationship between mortgage market rates and reversion rates. Historically, SVRs have been moving at a almost constant spread over market rates, but lenders stopped adjusting their SVRs downward as soon as the policy rate hit the zero-lower bound. Under considerable public and political pressure to pass-through the interest rate cut (The Guardian, 2008), lenders initially decided to lower SVRs. They did not however follow through once the BOE cut the base rate further by 150 basis points. Despite falling interest rates, average SVRs remained solid around 4%, and even increased in the following years.[4] As a consequence, reversion rates and rates on newly originated mortgages started to diverge and by the second half of 2013 the implied spread was higher than it had been before the crisis.

In the left panel in Figure 1.3, I plot the difference in annual payments between staying on the SVR and refinancing to a new 2-years fixed rate deal for a typical mortgage with a £100,000 remaining balance, 20 years left to pay down the principal and a LTV of maximum 75%. For around two years after the crisis, reversion rates were at the same level of market rates, or even cheaper. As financial markets recovered and risk premia went down, foregone savings for households that failed

[4] A similar pattern to the one observed in the UK is documented in Goggin et al. (2012) for the Irish market and the authors attribute the unwillingness of banks to pass-through interest rate cuts onto their reversion rates to increased market funding costs.

Figure 1.2: Mortgage Rates

The figure shows monthly time series of average quoted interest rates for different mortgage products and of the Bank of England base rate. The dark solid lines is the interest rate on a 2-year fixed initial period for a maximum loan-to-value of 75%. The lighter solid line is the corresponding rate for a 5 years initial duration. The dashed line is the average standard variable rate (SVR) applied by financial institutions. All mortgage rate series are taken from the Bank of England Interactive Database.

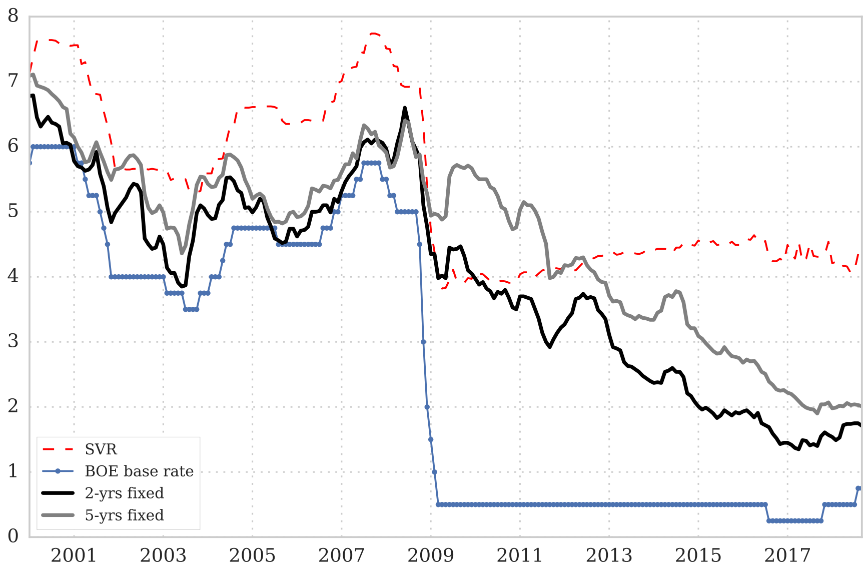

to refinance began to grow larger. Between 2013 and 2016, which is the period covered by my data, the average SVR was 4.44% against an average 2-years fixed rate of 2.18%, corresponding to an average annual difference in mortgage payments of £1,424.

The right panel of Figure 1.3 shows the change in mortgage payments implied by inaction relative to the expired deal. Right after the crisis, not only the incentive to refinance plotted in the left panel was small or negative, but borrowers reverting to the SVR would see their mortgage payments decrease substantially. From 2011, two things happen. First, the expected change in mortgage payments in case of inaction began to increase, and more and more borrowers reverting to the SVR would experience a jump in mortgage costs. Second, we observe a large difference in the experienced change in payments at reset for 2-years versus 5-years mortgages. On average, in the period 2013-2016, borrowers coming off a 2-years fixed deal saw their payments go up by £931 per year. At the same time, the rate reset meant a decrease in annual payments by £565 if the borrower had instead locked in the fixed rate five years before for five years. Because of steadily falling interest rates and high term premia after the crisis, borrowers on a 5-years contract were paying on average 2.3% points more in interest charges than borrowers on a 2-years contract resetting at the same time.

Standard models of optimal refinancing would predict no difference in refinancing behavior across borrowers that experience an increase versus a decrese in mortgage

Figure 1.3: Incentives to Refinance

This figure shows the evolution of the incentives to refinance over time, decomposed into the potential savings from refinancing (left) and the cost of staying on the SVR (right). Calculations consider a borrower with a repayment mortgage, £100,000 remaining balance and 20 years left until maturity. Specifically, the line in the left panel shows the change in annual payments when switching from the SVR to a new 2-years fixed rate initial deal, computed as $P(\text{SVR}_t) - P(r_t^{2yr})$. The right panel shows the change in annual payments when moving onto the SVR for mortgages coming off a 2-years (darker line) or a 5-years (lighter line) fixed deal, i.e it plots $P(\text{SVR}_t) - P(r_{t-2}^{2yr})$ and $P(\text{SVR}_t) - P(r_{t-5}^{5yr})$, respectively. At each point average quoted rates for mortgages with a maximum LTV of 75% are considered.

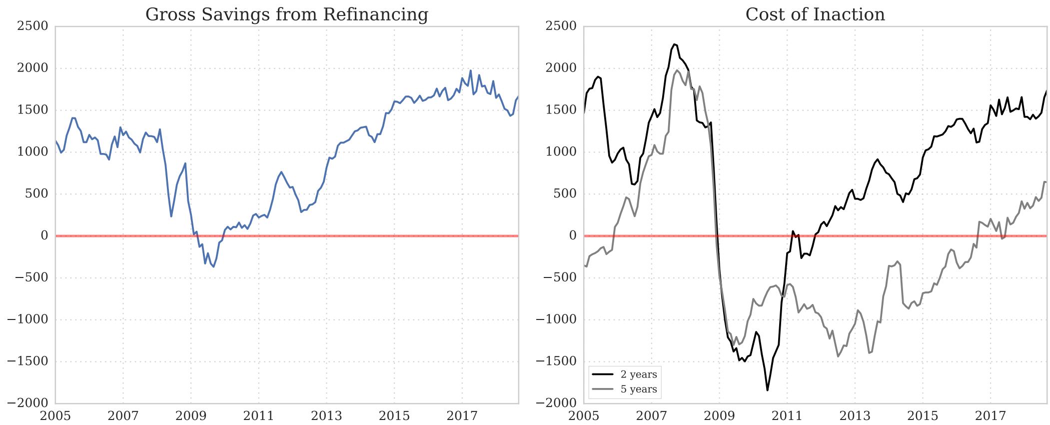

payments relative to the expired fixed rate. However, finding that this dimension matters in borrowers decisions, has implications for the effectiveness of monetary policy in reducing mortgage payments and stimulating household consumption. The pass-through of monetary policy is stronger on rates on newly originated mortgages due to competition among lenders (Scharfstein and Sunderam, 2016), thus the central bank has traction on refinancing incentives through the level of foregone savings. On the contrary, the cost of inaction, defined as the difference between current SVRs and past market rates, responds less and with a delay to monetary policy interventions. First, because it depends on the path of interest rates in the past and, second, because the pass-through of interest rates is limited by lenders' market power on their current clients.

1.3 The Theoretical Framework

Following the literature on optimal mortgage refinancing (e.g. Agarwal, Driscoll, and Laibson, 2013; Andersen et al., 2015), we can write the incentive to refinance as

$$I = r^0 - r^1 - x^* \tag{1.1}$$

where r^0 is the interest rate in case of inaction, r^1 is the interest rate on the new mortgage and x^* is a threshold level that captures the fixed cost of refinancing and the option value of refinancing in the future. All variables are time varying, but I drop the time subscript for convenience. A borrower that maximizes the utility in the outcome state will refinance if and only if the incentive to do so is positive, i.e. when $I > 0$. Agarwal, Driscoll, and Laibson, 2013 are the first to derive a closed form solution to the household's refinancing problem under a plausible set of assumptions, which other authors have used as rational benchmark against which to define refinancing mistakes (Keys, Pope, and Pope, 2016). Andersen et al., 2015 extend this rational model to incorporate inertia to generate heterogeneous responses to identical financial incentives. In particular, they allow the psychological cost that borrowers' associate with refinancing to vary across-borrowers, resulting in borrower-specific threshold levels x^*. The authors investigate how the estimated inertia covaries with borrower and mortgage characteristics and find it increasing in households' socio-economic status.

In this paper, I posit that borrowers have reference dependent preferences and evaluate the benefits from refinancing relative to individual reference points. Under this assumption, heterogeneous reactions to the same financial incentive may result from differences in reference points, even after controlling for borrower characteristics to proxy for financial literacy. This hypothesis is consistent with a model that assumes that household's utility function $u(C|R)$ depends on the consumption level C and a reference level of consumption R. Reference dependence is a fundamental principle in prospect theory and it is captured by the value function defined on the difference $C - R$ in Kahneman and Tversky, 1979. Building on this theoretical framework, I specify a borrower's utility function in terms of interest rates as

$$u(r|r^R) = \mu(r^R - r - \kappa(r^R, r)) \tag{1.2}$$

where r is a mortgage rate, r^R is the reference mortgage rate and $\kappa(r^R, r)$ is the potential cost involved with moving from the reference state to the new state, which reduces consumption in the outcome state. $\mu(\cdot)$ is a gain-loss function that satisfies the following properties:

A0. $\mu(x)$ is continuous for all x, twice differentiable and $\mu(0) = 0$

A1. $\mu(x)$ is strictly increasing.

A2. $\mu''(x) \leq 0$

Notice that I am not assuming a kink at the reference point, as in the well-known S-shaped value function in prospect theory. Because of the diminishing marginal utility resulting from concavity, the disutility from a loss is still larger in absolute value than an equally sized gain. What this specification does not assume, is diminishing sensitivity to losses, i.e. that the marginal disutility of a further loss in consumption decreases as the loss grows larger.

Figure 1.4: Value Function

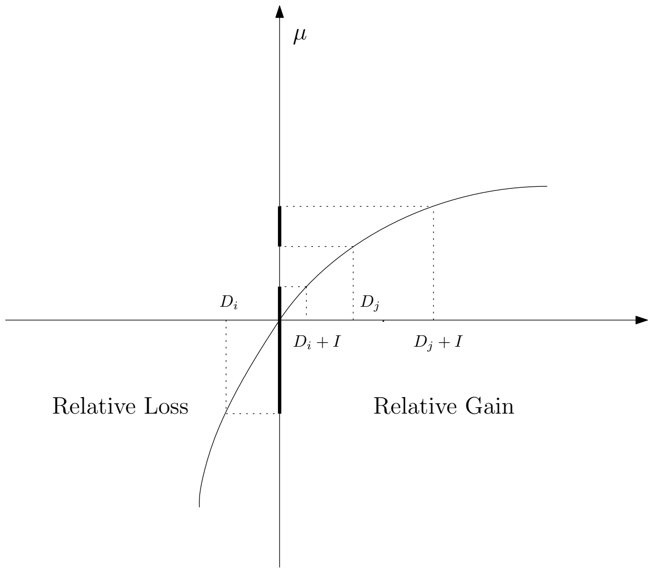

Assuming a strictly positive psychological cost of taking action ζ, a borrower i with reference rate r_i^R will refinance if and only if the change in utility is large enough to compensate for the hassle of refinancing, i.e.

$$\mu(r_i^R - r^1 - x^*) - \mu(r_i^R - r^0) = \mu(D_i + I) - \mu(D_i) > \zeta \tag{1.3}$$

where

$$D_i \equiv r_i^R - r^0 \tag{1.4}$$

is the distance of the reversion rate from the reference rate, expressed as a gain. I introduce the subscript i to stress that the reference point can vary across households for given r^0 and r^1. In the first term of the equation, $\kappa(r^R, r^1) = x^*$ captures the financial cost of refinancing that enters (1.1). By definition, there is no cost involved in reverting to the reversion rate r^0, so $\kappa(r^R, r^0) = 0$ in the second term.

Equation (1.3) implies that a positive rational incentive $I > 0$ is a necessary condition, but not a sufficient for $\zeta > 0$. Moreover, given concavity of μ, the propensity to refinance is negatively related to D_i. In other words, the higher the relative gain (or the smaller the relative loss) of reverting to r^0, the smaller the increase in utility from refinancing to a lower r^1. The intuition is visualized in Figure 1.4, where I draw an hypothetical value function to sketch the refinancing problem for two borrowers i and j with two different reference rates. D_i and D_j on the x-axis indicate the change in interest rates in case of inaction and determine the experienced change in monthly payments. D_i is negative meaning that $r_i^R < r^0$. In this case, inaction implies a loss in consumption and therefore a lower utility relative to the reference point. On the contrary, D_j is positive since $r_i^R > r^0$. Thus, even by doing nothing, borrower j experiences a drop in monthly payments once the mortgage rate resets. Given the same incentive I, refinancing implies an increase in utility indicated in bold on the y-axis for the two borrowers. However, notice that for borrower i, the perceived benefit from refinancing is much larger. This is because refinancing includes an additional increase in utility coming from avoiding a out-of-pocket loss relative to the reference point. The utility increase from gaining I, is instead much smaller for borrower j since, from her point of view, action only implies realizing an additional saving. Because of the assumption of decreasing sensitivity to gains, the same level of I might therefore not be enough to motivate borrowers who do not see the inaction state as a loss state into refinancing to a lower rate. The following testable hypothesis summarizes this idea:

Hypothesis (Reference Dependence)**.** *Given an action state and an inaction state, the incidence of refinancing in the cross-section of borrowers is negatively related to the difference D between the rate on the expired deal (reference rate) and the rate in the inaction state.*

1.4 Empirical Strategy

To test the hypothesis of reference dependence in refinancing decisions, I exploit the variation in reference points across borrowers whose fixed rate resets at a predetermined date. Formally, I run the following specification

$$Refinance_i = \alpha_{t,l,\mathrm{LTV}} + \beta D_i + \gamma' W_i + \varepsilon_i \tag{1.5}$$

where $Refinance_i$ is an indicator variable denoting whether loan i refinances after the reset and D_i is the distance between the reference rate and the reversion rate defined in equation (1.4). $\alpha_{t,l,\mathrm{LTV}}$ is a time*lender*LTV fixed effect and W_i is a vector of loan-level observables.

A key challenge to identify the effect of reference points on refinancing behavior is to control for differences in refinancing incentives. I leverage on three specific institutional features of the UK mortgage market to overcome this issue. In particular, I take advantage of the fact that (i) borrowers at the same lender face the same reversion rates at any point in time, (ii) mortgage pricing does not depend on borrowers' characteristics and (iii) borrowers can refinance to a different product with their current lender at a minimal hassle. As I explain in detail in the next paragraph, it follows that including granular three-way fixed effects for the time of the reset, the lender and the LTV of the mortgage at reset absorbs observable and unobservable heterogeneity in monetary incentives as defined in equation (1.1).

Rational incentives to refinance are positively correlated with the rate r_i^0 that the borrower is charged if she decides not to refinance. In the present context, this is the SVR to which the borrower automatically reverts at the end of the deal. This rate is constant for borrowers who have a loan with the same lender and whose initial rate resets at the same time, and it is observable. The second term r_i^1 in equation (1.1) is determined by the current level of interest rates and more precisely by the set of interest rates that are available to a given borrower at a given time. Mortgage products in the UK are very standardized and for any given product lenders offer the same interest rate to all borrowers that meet their lending standards. In particular, interest rates do not depend on individual borrowers' creditworthiness or other characteristics. Differently from the United States, where mortgage rates are quoted to borrowers individually and are a function of their credit score, default risk in the UK is priced based on the LTV ratio. Best et al., 2018 confirm in the data that after controlling for bank, time, interest rate (fixed or tracker), length of the initial deal and type of repayment (interest only or principal amortization), what determines the interest rates is the LTV ratio. In particular, the average mortgage interest rate is a step function of the LTV ratio, with sharp jumps (notches) at LTVs of 60%, 70%, 75%, 80% and 85%, and flat in between. Moreover, unlike in the US where many lenders offer different interest rates across states[5], in the UK mortgage rates do not vary across zipcodes or regions. Most products are available throughout the UK, even though some providers have limited lending areas. It is possible that some products are only available online, in branch or via an intermediary (MoneyFacts.com[6]). It follows that the same set of interest rates is in principle available to all borrowers that reset at a given time and whose LTV ratio falls within a given range, so that r_i^1 is constant within this cluster. The last term x_i^* is both increasing in the up front cost of refinancing and in the option value of waiting and refinancing at a future date. Refinancing is costly, both in terms of money and time. There is however a substantial difference between refinancing with the current lender (product transfer) or with a new lender. When transferring to a new mortgage with the current lender, no fees are usually charged and the procedure is commonly a matter of days and can be done entirely online. This is because for existing clients lenders usually do not require a new valuation of the property nor updated affordability checks, provided that the terms of the contract are unchanged. If the borrower wishes to modify the length of their term, increase the borrowed amount or change the repayment type of the loan, the lender will request a new assessment of both the financial situation and the value of the house. Thus, the cost of refinancing to the same product implies a different cost for new and existing clients of the mortgage provider. In turn, this affects refinancing incentives across borrowers, even though in principle they face the same set of available market rates. Lender fixed effects absorb this heterogeneity across borrowers at different lenders in accessing the same rate. Including lenders fixed effects also controls for differences in average refinancing probability, which may result from some lenders having more stringent requirements, higher fees, lenghtier procedures and different clienteles. Finally, the option value of waiting and to refinance in the future depends on the stochastic process of interest rates. I assume that, after including time fixed effects, borrowers expectations about future

[5]https://www.consumerfinance.gov/about-us/blog/7-factors-determine-your-mortgage-interest-rate/

[6]Moneyfacts is one of the most commonly used financial price comparison websites in the UK.

interest rates are unrelated to reference points.

The vector W_i includes control variables that are expected to influence borrowers incentive to refinance and that are not absorbed by the fixed effects. In particular, I control for the remaining balance on the mortgage given that, since interest savings from refinancing scale proportionately with mortgage size but refinancing cost is fixed, x^* is decreasing in mortgage size. Moreover, since the remaining time until maturity of the loan affects the option value of waiting, I include it as a control in the regression. Given the extensive set of fixed effects required to make robust inference, I first estimate a linear probability model.

At each point in time, the distribution of reference points across borrowers depends on the path of interest rates up to that point. Because of steadily declining interest rates, the right panel in Figure 1.3 shows that in the sample period D_i is positively related to the length of the fixation period of the maturing deal. The key identifying assumption for equation (1.5) to estimate the causal effect of D_i on refinancing decisions is that preferences for duration risk are uncorrelated with borrowers' propensity to show inertia or inattention or with other characteristics that may explain sluggishness in refinancing behavior. Borrowers' age and income have been shown in the literature to correlate with borrowers' responsiveness to refinancing incentives. Using data from the American Housing Survey (AHS), Campbell, 2006 shows that most active refinancers are younger, better educated, white households with higher-priced houses. Andersen et al., 2015 find similar results studying the Danish mortgage market. I can observe borrowers' age and income in the data, so I include them as controls in the regression. While age and income may affect the probability of refinance, they should not change the coefficient of interest since in the UK interest rates are not related to borrower characteristics. Still, since D_i is not exogenously assigned, I need to assume that there are no unobserved characteristics that are correlated with both D_i and refinancing decisions. In Section 1.7.1, I provide a set of additional results to rule out that my findings are driven by unobserved borrower characteristics simultanously driving duration and refinancing choices.

1.5 Data and Sample

1.5.1 Data Sources

For the main analysis I use a novel loan-level panel dataset on more than 2 million securitized residential prime mortgages in the UK provided by the European DataWarehouse (ED).[7] Data start in January 2013 and my sample ends in August

[7]ED collects loan-level information on the pool of loans backing RMBS that financial institutions pledge as collateral in Eurosystem refinancing operations. Following the eligibility requirements set by the ECB, since January 2013 participants to the Eurosystem have to submit updates on the underlying loans at least on a quarterly basis. As part of a measure to preserve collateral availability and market functioning, on September 6, 2012, the ECB extended eligibility to be used as collateral in Eurosystem credit operations to marketable debt instruments denominated in GBP (or US dollar or Japanese yen).

2017. In terms of coverage, the loans in the dataset correspond to roughly 10% of the total amount outstanding of mortgages in the UK over the period.[8] The dataset contains detailed information on the loans, including the name of the loan originator, the loan size, the origination date, the interest rate charged, whether the mortgage payments include amortization of the principal, the valuation of the property, the mortgage term over which the loan will be fully repaid, as well as the geographical location of the property at county level (NUTS3). The data also provides information on the purpose of the loan (purchase, re-mortgage, renovation, equity release, etc.), if it is a first or second mortgage and whether the mortgage is buy-to-let or owner-occupied. The data includes a number of borrower characteristics as of loan inception, namely age, income, income verification, employment status, credit score and whether the borrower is a first-time buyer.

The frequency of the data is either monthly or quarterly.[9] At each submission, I observe updated information on the payment history (current, in arrears, defaulted or prepaid) as well as the type of the mortgage and the interest rate charged. In particular, I know if the loan is currently on an initial (fixed or floating) rate, on the lender's SVR or on another rate (e.g. lifetime BoE base rate tracker, capped or discount). For introductory deals I observe the date when the deal ends and the loan reverts onto the follow-on rate unless the borrower refinances. In case the reset date is missing in the data, I recover it from the changes in the interest rate and the interest rate type across submissions. The dataset contains a variable that indicates the type of follow-on rate, whether it is the lender's SVR or another tracker rate.

Once the introductory period ends, I see from the following submissions if and when the borrower decides to refinance. If the borrower does not refinance, the loan appears on the lender's SVR. If the borrower decides to switch to a new product with the current lender, I observe the selected mortgage type and the new interest rate. Usually, borrowers that refinance after being on a fixed rate opt for a new initial fixed rate deal. If, instead, the borrower decides to remortgage with a different lender, refinancing appears in the data as a prepayment, after which the loan stops being reported.

The dataset contains updated information on loans' current balance and current loan-to-value ratio. For property values reported in the data and used to compute the current loan-to-value, I observe two reporting practices that vary across, but are consistent within, lenders. Most lenders update the value of the property at each submission according to an internal indexing methodology.[10] Other lenders,

Disclosure of loan-level information is also one of the eligibility requirements for credit operations with the Bank of England since November 2012. Since the Bank of England has access to the ED platform, many UK issuers use it to fulfill the disclosure requirements. Some issuers prefer alternative ways. Over the period 2013-2019, coverage of the UK RMBS market by the ED dataset varies between 30 and 60%.

[8]Data on total balances outstanding are from the Bank of England and the FCA.

[9]The frequency at which data are submitteed to ED depends on the RMBS coupon schedule.

[10]The data contain information on the valuation type used at each submission. Most of the properties are valued according to a full internal and external inspection at origination of the loan. On subsequent dates, if the property value is updated the valuation method is typically indicated as *Indexed*.

instead, report an updated property value only upon refinancing of the loan and usually only if the balance of the loan increases. Both behaviors are consistent with the fact that lenders usually do not require a new full valuation of the property in case of internal refinancing, except in case of modification of the loan terms. It is possibile that also lenders that report constant property values may use an indexing methodology before granding a product transfer to an existing client. For my analysis, I assume that the reported loan-to-value is the relevant one to determing the available set of refinancing rates. For robustness, I will include county fixed effects in some specifications in order to control for heterogeneity in house price growth across counties.

1.5.2 Sample Selection and Summary Statistics

Table 1.1 shows a snapshot of the cross-section of mortgage types as they are observed at the beginning of the sample in 2013. I see that a large fraction of borrowers (61%) are paying the SVR, which is surprising given that market rates in 2013 were already significantly cheaper. Since for borrowers with a small balance or few years left to maturity the potential savings might not justify the cost and hassle of refinancing, in the second row I restrict the sample to borrowers with a remaining balance higher than £100,000 and more than ten years to term. Since premia charged to highly leveraged borrowers increased substantially after the crisis, I also exclude mortgages with a LTV ratio higher than 75% to compute these figures. The fraction of mortgages on the SVR drops substantially to 43%, but it is still very large considering that borrowers in this subsample are foregoing substantial amount by failing to refinance to a new initial deal. To study reference dependence in refinancing decisions, I focus on mortgages on an initial fixed rate period.

In the early 2000s two major lenders guaranteed to their clients that their SVR would never rise more than 2% above the BoE base rate. Since the cut of the base rate in 2009, these reversion rates have been below most available market rates, which led the lenders in question to introduce a second, more expensive SVR for the newly originated mortgages. Since most mortgages in the dataset were originated before 2010, 30% of the loans in ED are on a low reversion rate as of their first submission. These loans are not considered in Table 1.1 since they imply low or negative incentive to refinance. Mortgages that are first observed on an initial fixed rate, but which are meant to revert to a low reversion rate, are also excluded from the analysis.

Like in most studies on the failure to refinance, one concern is that some borrowers might want to switch to a lower rate but cannot because ineligible to do so. The distribution of housing equity and unemployment might contribute to explain some of the observed sluggishness in refinancing behavior, since borrowers are likely to be denied refinancing if the account is in arrears, they have little or no equity, or there have been material changes in their circumstances (Agarwal et al., 2015b; Beraja et al., 2018). In order to exclude from the sample borrowers that may be unable to

Table 1.1: Distribution of Interest Rate Types (2013)

This table shows the fraction (in %) of mortgages by interest rate type at the beginning of the ED sample (2013). In the first row, reported figures are computed over all mortgages in the sample except for those on a lifetime tracker rate, which mostly indicates mortgages on a reversion rate that was guaranteed not to rise more than a small margin (usually 2%) above the BOE base rate. The second row only considers borrowers who have a remaining balance higher than £100,000, more than ten years left until maturity and a LTV ratio at reset not larger than 75%. The column SVR indicates the fraction of borrowers that are on their lender's Standard Variable Rate as of the first time they are observed in the data. Initial Fixed (Floating) indicates mortgages that are on an initial fixed (variable) rate and that will automatically revert to the SVR at the end of the deal. Other includes capped and discount mortgages.

	SVR	Initial Fixed	Initial Floating	Other	N
*All borrowers**	61.05	26.12	5.71	7.11	1'478'446
Borrowers with a strong incentive	43.18	42.13	8.07	6.62	227'424

qualify for a new mortgage due to bad performance or negative equity, I restrict my analysis to borrowers that were never reported late on their payments and who reset with a loan-to-value ratio below 90%. Still, these criteria cannot identify borrowers who are excluded from refinancing despite having positive equity and never missed a payment because they no longer meet lenders' eleigibility requirements, which have become stricter after the crisis and, in particular, since the introduction of the Mortgage Market Review in April 2014 in the UK. The cases of borrowers trapped in expensive reversion rates have attracted considerable attention in the popular press. According to those accounts, these so-called *mortgage prisoners* appear to be mostly self-employed or elderly people, who took out interest-only or self-certified mortgages. Moreover, as also reported in the FCA 2019 Mortgage Market Study, most of these mortgages are with unauthorized or inactive lenders, who do not offer any new deals. I do not expect this to be a major issue for my analysis for two reasons. First, there are no mortgages originated by currently inactive lenders in the data and servicer and originator are the same for the vast majority of loans. Second, in the regressions I control for income certification, repayment method, dummy for first-time buyers and age of the borrower, which I require to be non-missing. I also drop loans for which there is no information about the location of the property to rule out that my results are capturing heterogeneity in regional house price growth. In Section 1.6.3, I show that my results are robust to these restrictions.

The final sample constains 85,830 reset events distributed fairly homogeneously between January 2013 and August 2017. Table 1.2 presents summary statistics for the mortgages in the sample. The first four columns show number of observations, mean, standard deviation and median for loans that reset to a SVR that is lower than the rate on their past fixed deal. The next four columns present the same summary statistics for borrowers that instead reset to a SVR that is higher than their expired fixed rate. 35% of the mortgages in the sample reset with $D_i > 0$, i.e. their

Table 1.2: Summary Statistics for Mortgages that Experience a Rate Reset

This table shows summary statistics for the mortgages in the analysis sample, which includes only mortgages that see their initial fixed rate reset at some pre-determined date during the sample period. I report number of observations, mean, standard deviation and median of the control variables separately for mortgages that experience an automatic drop in monthly payments at reset (GAIN) and those who see their payments increase when falling on the reversion rate (LOSS). All variables in Panel A and B are measured at the end of the introductory deal, except for gross income and credit score which are as of the inception of the loan. Loan-to-income is current balance divided by income at inception. Income verification indicates whether the income of the borrower has been verified by the lender when the loan was granted. Reported credit scores in the ED data follow different scales depending on the score provider: Callcredit, Experian or Equifax. To allow comparability, I standardize the score by the maximum value in each score system. For lenders that assign credit scores according to internal systems, I standardize using the maximum assigned value in sample. Panel C reports original loan amount, LTV ratio and term of the loan (in years) as of the origination of the mortgage.

	Reset implies GAIN ($D_i > 0$)				Reset implies LOSS ($D_i < 0$)			
	N	Mean	SD	Med	N	Mean	SD	Med
Panel A: Loan Characteristics at Reset of the Current Initial Fixed Deal								
Interest Rate (%)	30336	5.32	0.83	5.19	55494	3.15	0.58	3.24
Loan Value (1000)	30336	93.62	66.52	81.18	55494	105.24	92.27	81.86
Loan-to-Value (%)	30336	59.32	25.67	67.47	55494	43.84	19.64	44.61
Years to Maturity	30336	17.49	7.85	17.92	55494	13.92	6.62	13.67
Interest Only	30336	0.10			55494	0.22		
Panel B: Borrower Characteristics								
First-Time-Buyer	30336	0.41			55494	0.19		
Gross Income (1000)	30336	38.78	23.50	33.32	55494	46.64	31.59	38.00
Age at Reset (Years)	30336	43.21	10.88	42.00	55494	47.46	9.62	47.00
Loan-to-Income	30336	2.51	1.17	2.48	55494	2.34	1.28	2.21
Income Verification	30336	0.62			55494	0.52		
Credit Score	15475	0.75	0.12	0.75	23371	0.78	0.14	0.77
Panel C: Loan Characteristics at Origination								
Original Loan Value	30336	106.81	68.39	93.00	55494	126.61	95.83	102.00
Original Loan-to-Value	30331	70.24	23.38	78.40	55481	59.99	21.45	63.60
Original Loan Term (Years)	30336	23.67	6.38	25.00	55494	20.99	6.04	22.00

mortgage payments are lower on the SVR than on the introductory rate. Figure 1.9 in the Appendix shows the time series of SVRs by lender.

1.5.3 Distance from the Reference Point

The main explanatory variable of interest is the distance D_i between the reversion rate and the reference rate, defined in equation (1.4). The left panel of Figure 1.5 plots the distribution of this measure in the data. Recall that D_i is the difference in percentage points between the interest rate charged on the introductory deal and the SVR charged in case of inaction. A negative value thus means that reverting to the SVR implies an increase in mortgage payments relative to what the borrower has

Figure 1.5: Variation in D_i

The left panel plots the distribution of D_i defined in equation (1.4) as the percentage point difference between the rate paid on the introductory period and the reversion rate (SVR). The right panel shows the distribution of the change in annual mortgage payments implied by D_i (in £).

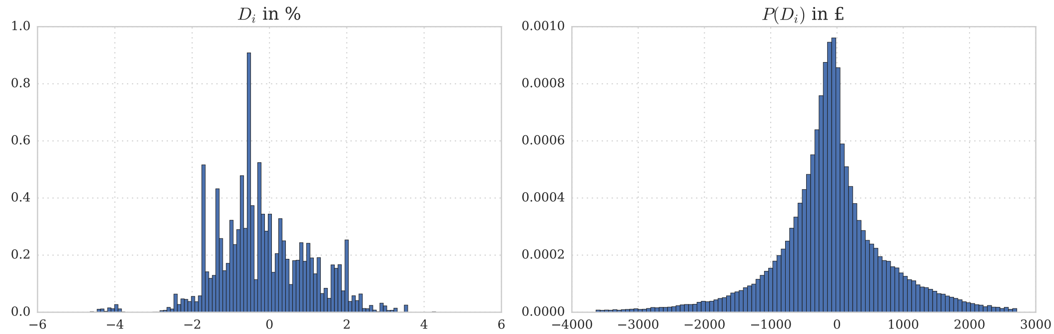

been paying until that moment. In other words, staying on the SVR implies a loss in disposable income from the perspective of the reference point. In my sample, about 40% of resets happen when this distance is negative. The average distance is 9.5 basis points. Since interest rates have been going down over the period under consideration, the fraction of borrowers with a negative distance decreases over time, dropping from 53% in 2013 to 32% in 2017. In the right panel, I plot the change in annual disposable income in the inaction state relative to the reference state in the right panel for the borrowers in the sample, which are roughly normally distributed between minus and plus £2,000.

1.6 Results

1.6.1 Graphical Evidence: Losses versus Gains

In this section, I provide preliminary evidence that reference points matter for borrowers' refinancing decisions. To do so, I group households based on whether the SVR is higher (LOSS) or lower (GAIN) than the rate on the fixed deal. Figure 1.6 plots the fraction of households that have refinanced their mortgage within six months from the end of the introductory fixed deal, by year of the reset. The left panel considers mortgages with a LTV at reset between 20% and 60%, the center panel mortgages with a LTV between 60% and 75% and the right panel mortgages with a LTV between 75% and 85%. The average refinancing rate across observations is 62.5%.[11] The reference dependence hypothesis posits that we should observe a

[11] Using a comprehensive dataset on the universe of UK mortgages, the Financial Conduct Authority (FCA) finds that in the period 2015-2016 more than 75% of the mortgages have been refinanced within six months from the reset. For the same period, the refinancing fraction in my sample is lower, 65.8%. The FCA 2019 Mortgage Market Study can be found at https://www.fca.org.uk/publication/market-studies/ms16-2-3-final-report.pdf.

Figure 1.6: Reference Points and Refinancing Activity

The figure shows the fraction of households (with at least five years left until maturity) that refinance within six months from the end of the introductory fixed deal by year and LTV class. The color of the bar indicates whether inaction implies a loss (dark grey) or a gain (light grey) relative to the state before the rate reset.

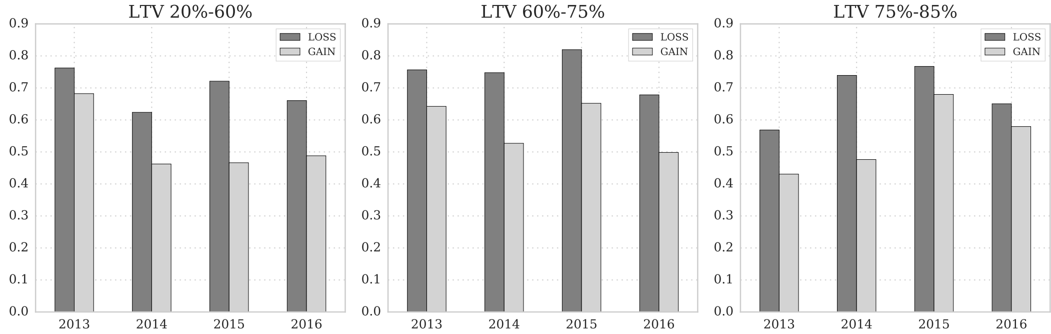

higher refinancing rate among households for which inaction implies a loss in disposable income relative to the pre-reset period. Consistent with this prediction, the average refinancing rate among this group of borrowers is 18.5% points higher than for mortgages with a maturing fix rate lower than the reversion rate. Households for which the inaction state implies a loss refinance more in each year under consideration and in each LTV group. For 40% of the loans I can observe their performance over two years after they reset. The difference between the two groups is still large and significant (14.0%).

1.6.2 Regression Results

Table 1.3 presents estimation results of the model specified in equation (1.5). The dependent variable $Refinance_i$ is set equal to one if borrower i has refinanced within six months from the end of the introductory period.[12] The first regression demonstrates that the negative relationship between reference points and refinancing decisions observed in Figure 1.6 is statistically significant. The coefficient of -0.066 (t-stat = -5.76) indicates that an increase by 1 percentage point of the difference D between the reference rate and the reversion rate is associated with a 6.6 percentage point decrease in the individual's probability of refinancing after the initial fixed rate period ends. The unconditional probability of refinancing within six months from the rate reset is 59.1% and the standard deviation of D is 1.25 percentage points in the sample. Therefore, a 1-standard deviation increase in the perceived gain from moving to the reversion rate decreases the probability of refinancing by 14% of its unconditional value.

From column (2) all specifications include reset month $\times$ lender $\times$ LTVbin fixed effects to control for unobserved heterogeneity in financial incentives to refinance

[12] I show the robustness of the results to alternative choices of the refinancing horizon in Section 1.6.3.

Table 1.3: Reference Point Effects on the Decision to Refinance

The table shows OLS regressions of different variants of equation (1.5). The dependent variable is a 0/1 indicator for whether the mortgage is refinanced within six months from the end of the initial fixed period. The explanatory variable of interest is the distance from the reference rate D, defined for each individual borrower facing a rate reset as $D_i = r_i^R - SVR_i$ (in percentage points). r_i^R is the expired fixed rate of mortgage i and SVR_i is the corresponding reversion rate, i.e. the SVR of loan's lender l. Reported t-statistics in parentheses are clustered at the month when the rate resets and at the region (nuts-2) where the property is located. Key: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
D	-0.066***	-0.079***	-0.063***	-0.062***	-0.062***	-0.059***	-0.066***	
	(-5.771)	(-13.679)	(-9.860)	(-9.736)	(-9.812)	(-9.252)	(-10.346)	
GAIN Dummy ($D > 0$)								-0.134***
								(-10.011)
Controls								
Log Balance			0.077***	0.100***	0.115***	0.112***	0.144***	0.150***
			(7.248)	(9.567)	(16.731)	(16.454)	(17.159)	(18.994)
Years to Maturity			0.006***	0.003***	0.003***	0.003***	-0.000	-0.000
			(6.946)	(3.384)	(3.444)	(3.194)	(-0.375)	(-0.298)
LTV			0.002***	0.001**	0.001	0.000	0.001*	0.001
			(3.502)	(2.175)	(1.090)	(0.760)	(1.883)	(1.513)
Borrower's Age				-0.003***	-0.003***	-0.003***	-0.003***	-0.003***
				(-4.048)	(-4.529)	(-4.621)	(-3.859)	(-4.305)
Log Income				-0.023***	-0.024***	-0.025***	-0.047***	-0.047***
				(-4.173)	(-4.206)	(-4.341)	(-7.178)	(-7.379)
Repayment Mortgage							0.156***	0.155***
							(11.523)	(11.486)
Income Verified							-0.036***	-0.037***
							(-7.303)	(-6.853)
Firsttime Borrowers							-0.034**	-0.037***
							(-2.498)	(-2.760)
Month * Lender * LTV bin FE		✓	✓	✓	✓	✓	✓	✓
County FE					✓	✓		
Origination Year FE						✓		
Observations	79,468	79,468	79,468	79,468	79,467	79,467	79,468	79,468
R-squared	0.028	0.126	0.148	0.151	0.155	0.156	0.162	0.160

across borrowers. As I explain in detail in Section 1.4, given mortgage pricing in the UK available refinancing rates, refinanging costs and eligibility criteria are assumed to be constant within borrowers at a given lender with a similar LTV whose initial deal ends at the same time. The coefficient on D becomes more negative, dropping to $-.079$ (t-stat = -13.69). In the third specification I control for the log of the outstanding loan balance, the remaining years to maturitiy, the age of the borrower, the log of the household's income and the LTV. Notice that outstanding balance, years to maturity and LTV are all measured prior to the reset of the rate to make sure that they do not capture potential ex-post decisions to increase the borrowed amount or to extend the maturity of the loan upon refinancing. The inclusion of the controls reduces the estimated effect of reference points on refinancing choices slightly. The coefficient on D changes to $-.062$ (t-stat = -9.75) indicating that differences in mortgage and borrower characteristics also help explain differences in refinancing behavior. I verify that most of the change is driven by the inclusion of log balance, while age and income leave the coefficient essentially unchanged in column (4).This indicates that reference points do not covary significantly with demographics, which mitigates the concern that the results might be explained by unobserved borrower heterogeneity. Column (5) additionally controls for county fixed effects and column

(6) for origination year fixed effects. Both the sign and the significance of the coefficient on D are unaffected across columns, indicating that the results are not driven by a particular combination of fixed effects. In column (7) I include a set of dummies for repayment mortgage, income verification and first-time buyer. Controlling for these factors might be potentially important since they may both affect the observed rate on the mortgage and the probability that a borrower is able to refinance. In particular, individuals whose income was not been verified during the mortgage application are usually self-employed people who self-certified their income and who are therefore more likely to be denied refinancing. Including them in the regression makes the result, if anything, stronger. Finally, column (8) estimates the regression on a binary variable GAIN that takes the value of one when $D_i > 0$, and zero otherwise. The coefficient indicates that the average refinancing probability of borrowers whose mortgage payments are expected to go down regardless of their actions at the end of the fixed rate period is 13.4 percentage points lower than that of borrowers who would experience an actual loss in disposable income if they did not refinance.

Table 1.4 examines this mechanism further by estimating the model across subsamples based on quintiles of borrower age, borrower income and loan balance. The literature finds that refinancing mistakes are related to proxies for financial literacy and are more common among older borrowers and households with a lower income. Estimating the relationship separately within age and income groups allows me to test whether reference point effects weaken as borrowers are expected to be more financially savvy. There is a large debate about whether individuals make behavioral mistakes only when the financial consequences are negligible (Agarwal et al., 2015a; Pope and Schweitzer, 2011). Since potential savings are higher for borrowers with a larger outstanding debt, I check if behavioral biases disappear when the stakes are high. Before looking at the regression results, notice that the fraction of refinancers reported in the last row of each panel shows patterns consistent with previous findings in the literature: The probability to refinance is decreasing in age and increasing in both income and outstanding balance. Turning to the coefficient on the GAIN dummy, we see that overall the Table show that the heterogeneity in refinancing behavior related to differences in reference points exists over and above any effect from demographics. The coefficient is always negative and significant, and shows little variation across quintiles of the covariates. The last column in each panel reports the estimated differential effect of going from the bottom quintile to the top quintile of the distribution of the relevant grouping variable. In particular, the table shows the estimate and the respective t-statistic of the coefficient on GAIN $\times$ Q5 from a regression of $Refinance_i$ on levels and interactions of GAIN dummy and quintile dummies. While young mortgagors appear to be less affected by reference points and the difference between the first and last age quintile is statistically significant, the effect is not linear across quintiles. The effect of the reference mortgage rate is instead unchanged across income groups, while it decreases with the size of the outstanding balance.

So far, the analysis has relied on within lender, time and LTV bucket variation to control for potential omitted variable bias coming from heterogeneity in savings from refinancing. An alternative way of doing this would be to include on the right-hand side of the regression both the relative gain D and the actual average saving, namely the difference between the SVR and current market rates. I do this in Table 1.5. The advantage of this specification compared to the fixed-effect regression is that it allows to estimate the sensitivity of borrowers to both actual and relative gains. The disadvantage is that incentives to refinance are not determined solely by the interest rate saving but also depend on expectations about future changes in interest rates. These expectations are likely to change over time and depend on macro events and policy announcements.

From the specification in column (1), I estimate that a 1% increase in the relative gain decreases *ceteris paribus* the probability of refinancing by 5.1 percentage points. This is lower than the 6.6 percentage points estimated from the within groups regression in Table 1.3, but still negative and strongly significant. The coefficient on $SVR_i - \overline{r^{2yr}}$ indicates that a 1% increase in the savings from switching to a lower market rate increases the probability to refinance by 9.6 percentage points. The relative magnitude of the coefficients shows that the impact of reference points is roughly half as large as the impact of rational incentives. Said differently, these estimates imply that in order to produce the same effect on refinancing, the stimulus coming from lower interest rates has to be 1.5 times stronger in the presence of reference dependence, than without it. Column (2) includes additional controls for the type of mortgage, and the result becomes even stronger. In column (3), I control for the savings coming from refinancing to a 5-years fixed rate mortgage. The coefficient is still positive, but is not statistically significance. Columns (4) and (5) provide a robustness test for the results in columns (2) and (3), respectively. Instead of using the actual SVR that applies to individual loans, I use the average SVR at that time. The coefficient on the saving from switching in column (4) is slightly smaller and losses some significance. All specificantions include lender and LTV buckets fixed-effects.

Table 1.4: Effect by Quintile of Age, Income and Loan Balance

The table reports estimation results of specification (8) in Table 1.3 across different subsamples. The dependent variable is a 0/1 indicator for whether the mortgage was refinanced within six months from the end of the initial fixed period. GAIN dummy is equal to one when $D_i = r_i^R - SVR_i > 0$, which indicates that borrower i will experience a drop in mortgage payments if she does not refinance and stays on the lender's SVR. The table shows how the coefficient on the GAIN dummy changes across age quintiles (Panel A), income quintiles (Panel B) and loan value quintiles (Panel C). Included control variables in each panel are log balance, years until maturity, age of the borrower, log income, LTV, repayment method dummy, income verification dummy and a dummy for whether the borrower is a first-time buyer. In each panel, the second to last row indicates the average value of the grouping variable in each quintile and while the last row reports the respective refinancing frequency. In the last column, GAIN*ΔQ reports the coefficient on the interaction term of GAIN dummy and Q5 estimated from a regression with interaction terms between GAIN dummy and quintile dummies and where the base quintile is Q1. Reported t-statistics in parentheses are clustered at the month when the rate resets and at the region (nuts-2) where the property is located. Key: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

	Low Q	Q2	Q3	Q4	High Q	GAIN*ΔQ
Panel A: Age Quintiles						
GAIN Dummy	-0.087***	-0.140***	-0.146***	-0.145***	-0.135***	-0.071***
	(-4.981)	(-10.955)	(-9.632)	(-7.529)	(-7.774)	(-4.746)
Observations	15,773	16,542	18,202	13,371	15,426	79,477
R-squared	0.161	0.158	0.153	0.152	0.187	0.163
Age (mean)	32.0	40.2	46.5	51.9	60.8	
Refi Fraction (%)	61.1	63.3	62.3	59.7	47.2	
Panel B: Income Quintiles						
GAIN Dummy	-0.131***	-0.139***	-0.137***	-0.131***	-0.131***	-0.003
	(-11.071)	(-7.367)	(-7.409)	(-6.724)	(-8.131)	(-0.196)
Observations	15,882	15,764	15,762	15,846	16,070	79,477
R-squared	0.168	0.165	0.167	0.158	0.170	0.161
Income (mean in £1000)	17.9	27.6	36.3	48.3	89.4	
Refi Fraction (%)	50.7	57.3	60.1	62.5	63.9	
Panel C : Loan Balance Quintiles						
GAIN Dummy	-0.143***	-0.147***	-0.127***	-0.118***	-0.123***	0.027**
	(-11.579)	(-7.395)	(-8.354)	(-6.548)	(-8.271)	(2.039)
Observations	15,572	15,893	15,714	15,934	16,208	79,477
R-squared	0.122	0.119	0.145	0.162	0.182	0.162
Loan Balance (mean in £1000)	27.2	56.0	82.0	116.0	224.0	
Refi Fraction (%)	41.3	58.2	62.3	65.1	67.1	

Table 1.5: Horse Race Regression: Actual versus Relative Gains

The dependent variable is a 0/1 indicator for whether the mortgage is refinanced within six months from the end of the initial fixed period. The explanatory variable of interest is the distance from the reference rate D, defined for each individual borrower facing a rate reset as $D_i = r_i^R - SVR_i$ (in percentage points). r_i^R is the expired fixed rate of mortgage i and SVR_i is the corresponding reversion rate, i.e. the SVR of loan's lender l. $\overline{r^{2yr}}$ is the average rate on a newly originated 2-years initial fixed deal at the time of the reset. Similarly, $\overline{r^{5yr}}$ is the rate on a newly originated 5-years initial fixed deal at the same time. $\overline{SVR}$ is the average quoted SVR across all lenders. The first set of control variables includes log balance, years until maturity, age of the borrower, log income and current LTV. The second set of controls corresponds to a repayment method dummy, an income verification dummy and a dummy for whether the borrower is a first-time buyer. Reported t-statistics in parentheses are clustered at the month when the rate resets and at the region (nuts-2) where the property is located. Key: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

	(1)	(2)	(3)
D_i	-0.051***	-0.054***	-0.066***
	(-6.445)	(-6.788)	(-10.717)
$SVR_i - \overline{r^{2yr}}$	0.096**	0.097**	0.049***
	(2.039)	(2.040)	(4.481)
$SVR_i - \overline{r^{5yr}}(\perp)$			0.040**
			(2.102)
Controls	Yes	Yes	Yes
Lender FE	Yes	Yes	Yes
LTV Bucket FE	Yes	Yes	Yes
Observations	79,468	79,468	79,468
R-squared	0.113	0.124	0.129

1.6.3 Robustness: Alternative Horizons, Sample and Estimation Approach

I start by ensuring that the results presented in Table 1.3 are robust to alternative refinancing horizons. For the main analysis, I define the dependent variable to be equal to one if the loan is refinanced in the 6-months period from the reset date. Table 1.11 repeats the previous analysis for different event windows. The first column defines refinancing over a 3-months period, while the third and fourth columns consider refinancing within 9 and 12 months, respectively. For convenience, in the second column I report again the results for the 6-months period. Results are robust to all choices of horizon. Moreover, the analysis shows that the effect of the distance from the reference point is persistent and even becomes stronger over time. The last row of the table reports the fraction of borrowers who refinance and shows that most borrowers that decide to refinance do so within the first 3 months from the reset date. The refinancing fraction increases by an additional 10% over the following 9 months, indicating that some borrowers might need more time before becoming active, possibly because they were not paying attention and did not immediately realize that the interest rate has changed and that refinancing might be optimal. Taken together, these results show that refinancing in later months mostly comes from borrowers who see their mortgage payments increase automatically, rather than from those borrowers who face a gain and whose refinancing activity is catching up.

Recall from Section 1.5 that to be included in the analysis, I require loans to have non-missing information about several borrower and contract characteristics in order to address a potential omitted variable bias. I relax this restriction in Table 1.12, increasing the sample size by 50%. Specifically, I allow for age of the borrower, location of the property, repayment type and income verification to be missing. For the sake of comparison I repeat the analysis in the restricted sample, dropping age as a control and without including regional fixed effects. The main results are robust to this increased sample.

In the main analysis, I rely on an OLS framework to estimate the effect of reference points on the refinancing decision. This is done to ensure robust inference by including a large number of fixed effects. Table 1.13 presents results from a Logit model and confirms the previous finding that reference points are an important determinant of the refinancing outcome. The Table reports average marginal effects and indicates that a one percent increase in the distance from the reference point D is associated with a decrease in the probability of observing a refinancing within six months between 6.5% and 7.9%.

1.7 Addressing Identification Concerns

To test for reference point effects on refinancing decisions, ideally one would randomly assign reset dates and hence reference rates to borrowers. The empirical analysis in the previous section exploits cross-sectional variation in reference points that comes from borrowers' choices about the fixation period. Given steadily falling interest rates, the sample period contains little variation in the distance from the reference point that is orthogonal to the length of the fixation period. In particular, in 2013-2017 there is a positive correlation between the reference point and the length of the fixation period, as it is clear from the plot in the right panel of Figure 1.3. In this section, I address three challenges to my identification strategy that arise from the lack of exogenous variation in reference points.

First, one might be concerned that, despite controlling for borrowers' age and income, the results are explained by unobserved differences in borrowers' propensity to refinance. For example, if borrowers that prefer longer fixation periods are also more likely to be inattentive at reset of the fixed rate or they attach a higher psychological cost to taking action, this would lead to a negative relationship between reference rates and the probability to refinance.

The second identification challenge comes from the fact that mortgagors in the UK usually refinance several times before the principal is paid down: At the end of each introductory period most borrowers switch to a new initial deal so that mortgages basically turn into a succession of short-term contracts. Clearly, a borrower choosing a succession of 5-years contracts will be faced with the decision to refinance less frequently than if she had chosen 2-years contracts. This raises two concerns for my analysis, namely survivorship bias and learning effects from frequent refinancing. Consider two borrowers who took out a mortgage roughly at the same time and whose current initial fixed period expires on the same date. Assume that one of them was on a 5-years contract while the other one on a 2-years contract. If borrowers have time-invariant preferences for a given duration, we expect the borrower with the shorter deal to have refinanced more frequently in the past. On the one hand, this means that there were more occasions in which she could have failed to refinance even when it was optimal to do so. The fact that I observe the borrower in the sample means that she did refinance before, which might introduce selection bias. Moreover, we know from the household finance literature that borrowers appear to learn from repeated financial decision making (Agarwal, Rosen, and Yao, 2016). Another related concern might therefore be that borrowers with shorter fixation deals are faster at reacting to refinancing incentives because they have more experience with the refinancing process, which reduces the hassle of taking action.

Finally, one may worry that borrowers that select into longer maturities have a higher probability of being denied refinancing. This is a concern given that Hertzberg, Liberman, and Paravisini, 2018 show from the peer-to-peer lending market in the US that borrowers with higher unobservable repayment risk tend to self-select into longer maturity contract based on private information. As a consequence, lenders may use the maturity of the previous loan as a device to screen observationally identical borrowers. While I exclude borrowers that are late on their payments or who are close to negative equity, and I control for mortgage types that are more likely to be refused refinancing, I cannot exclude this channel a priori. To control for the incentive to self-select based on unobservables, I control for the term-premium at the time of the origination of the mortgage. The idea is that the larger the difference in term premia between 5-years and 2-years maturities, the worse the adverse selection problem gets. I show that the impact of the relative gain remains strong and significant after including controls for the incentive to self-select.

In Sections 1.7.1 to 1.7.3, I present evidence that mitigates the concern that these alternative explanations might be driving the results.

1.7.1 Endogenous Fixation Period

The first identification concern is that borrowers who prefer higher duration risk are also *ceteris paribus less likely* to refinance at the end of the fixed deal. Theory predicts that households with a large mortgage, uncertain labor income, high risk aversion, high cost of default, and low probability of moving should choose to bear less interest rate risk (Campbell and Cocco, 2003). If anything, risk averse borrowers should be relatively *more* careful not to fall on the variable rate of the lender, and households that do not plan to move in the near future should find the flexibility of the SVR *less* attractive. Admittedly, however, the household finance literature on mortgage choice finds that borrower's age, which is often associated with sluggishness in refinancing behavior, also suggests a preference for fixed rate mortgages over variable rate ones.[13] Overall, this assumption is inherently untestable and given the data I cannot rule out with certainty its effect on my results. Nevertheless, the results presented in the following sections show that a preference for longer fixation periods are not generally associated with lower propensities to refinance.

In Section 1.7.1 I first show that at least in terms of observable characteristics, namely age and income, people that choose two versus five years contracts are largely similar. Then, in Sections 1.7.1 and 1.7.1, I show using two distinct approaches that, in the absence of differences in reference points, a higher fixation period does not seem to be associated with a lower probability to refinance.

Mortgage Duration Choice and Borrower Characteristics

There is substantial evidence that older and lower income borrowers tend to react less to refinancing incentives (Bajo and Barbi, 2018; Andersen et al., 2015; Keys, Pope, and Pope, 2016). A first indication that differences in unobserved characteristics are unlikely to explain my findings is the fact that the estimated coefficient on D_i in Table 1.3 is virtually unchanged when I include age and income in the regression. If the choice of mortgage duration was correlated with the probability that a borrower is subject to behavioral biases such as inattention and inertia, we would expect to observe a strong correlation between reference rates and these demographics.

In this section, I compare observable characteristics of borrowers that take out a mortgage in the same month, but choose different fixation periods. There are several data limitations that I have to overcome for this purpose. First, while I know the end date of the introductory deal, the data do not contain information on loans fixation period. I can however recover the fixation period from the time series of submissions, provided that the loan has enough observations prior to the event or was originated not more than two years before the first submission to ED. Second, to answer the question at hand I would need to observe income as of when the choice about the fixation period was made. However, income in the database is at inception of the loan. The longer the time gap between loan origination and the latest fixation choice, the less accurate this information is going to be. I therefore restrict the analysis to loans that were originated in 2011, for which I know that the first observed reset corresponds to the first reset since origination and that the reported income is the income at the moment of the decision about duration. The sample contains roughly 6,000 observations. 75% of the loans have an fixation period of two years, 17% of five years and the rest is split between three and four years. This dominance of short fixation contracts is consistent with Koijen, Van Hemert, and Van Nieuwerburgh, 2009 who find that when term premia are high households tend to bear more interest rate risk.

[13] Paiella and Pozzolo, 2007 find using Italian survey data that household head's age is negatively related to the probability of taking out an ARM versus a FRM.

Table 1.6: Age and Income by Fixation Period

This table shows number of observations, mean, standard deviation as well as the 25th, 50th and 75th percentile of the distribution of age and income of borrowers that chose either a 2-years or a 5-years initial deal. The number of observations is much smaller than in the full sample because the ED data do not explicitely report the fixation period of each loan. Whenever possible, the length of the fixation period is recovered from the time-series of the submissions, as I explain in detail in section 1.7.1. Moreover, to make sure that income reflects borrower income at the time of the choice about the fixation period, I only consider mortgages that were originated in 2011. The last column reports the estimated coefficient on the 5-years indicator in the regression $Characteristic_i = \alpha_t + \alpha_k + \delta 5yrs_i + \gamma Controls_i + \varepsilon$ for loan i originated in month t in region k. For this regression, I include Loan-to-Value, Loan-to-Income and loan term as control variables. t-statistics are in parenthesis. Standard errors are clustered at the region level.

	2 years						5 years						δ
	N	Mean	SD	25%	50%	75%	N	Mean	SD	25%	50%	75%	
Age (Years)	4266	34.7	8.7	28.0	33.0	40.0	999	35.6	9.4	28.0	33.0	43.0	2.409*** (15.450)
Income (£1000)	4266	43.1	26.9	26.3	36.2	50.8	999	41.0	28.2	24.2	33.5	47.7	-1.859* (-1.917)

Table 1.6 shows the distribution of borrowers' age and income separately for 2-years and 5-years initial deals. Despite a much smaller number of observations for the 5-year deal, borrowers look very similar across the two groups. The last column shows the average differences and associated t-statistics for borrowers that choose a 5-year deal relative to those that choose a 2-year deal, estimated in a regression where I additionally control for mortgage characteristics as well as month and region fixed-effects. I find that borrowers with that choose the longer fixation period are 2.4 years older and have a gross annual income that is £1,860 lower on average. Even though the coefficients are statistically significant (for income only at the 10% level) and the sign indicates that this channel might amplify my results, the differences are small in magnitude confirming the result from the unconditional distribution.

Evidence from the 2007-2008 NMG Survey

The evidence presented in the previous paragraph that borrowers are similar across fixation periods is reassuring. However, the concern remains that observationally identical borrowers who self-select into different contract lengths differ along unobservable and hard-to-measure dimensions that explain different refinancing behavior. Ideally, we would like our sample to contain periods when reference rates are the same or lower for 5-years deals and 2-years deals. In this way, one would be able to directly test whether differences in fixation periods may be amplifying the effect of reference points. Observing that borrowers who choose to fix the interest rate for longer tend to refinance just as often or more often than the average borrower would provide support for our assumption, and strengthen the evidence in favour of the reference dependence hypothesis.

The right panel in Figure 1.3 shows that in the years 2007-2008 fixed rates were almost identical for borrowers coming off a two or five years introductory deal. Since these years are not in the ED database, I use data from the *NMG Survey of Household Finances* that the BoE carries out on an annual basis since 2004 to get information about refinancing activity during that period. At the end of 2007 and 2008, 908 respondents with a fixed rate mortgage on their house were asked whether their initial deal had expired within the previous twelve months. Out of these, 219 households did report a positive answer. For these households, I can infer whether the mortgage was refinanced from the type of interest rate the household reports to be currently paying (fixed, variable, SVR, etc.). Because the survey does not provide information on the length of the initial period nor the interest rate paid on the recently ended deal, I use the presence of dependent children (aged below 15) in the household as a proxy for having a longer initial fixation period. While the literature on mortgage choice finds little explanatory power of borrowers characteristics in predicting the preference for ARMs versus FRMs, a notable exception next to the age of the borrower is that families with children tend to insure more against interest rate risk (Paiella and Pozzolo, 2007). More generally, there is evidence in the literature that households with children tend to be more risk averse. The survey also contains information about the age of the respondent, the remaining balance on the mortgage, the current value of the property[14], total (secured and unsecured) outstanding debt, income and level of education. Figure 1.7 in the Appendix plots income and age distributions of borrowers in the survey and in the ED loan-level data. Average annual household income in ED is £43,877, but higher among survey respondents (£49,750). Homeowners in the survey are also younger, with an average age of 42 years compared to 46 years in ED.

In Table 1.7 I report estimated difference in average refinancing probability of families with children relative to families without children. I find that the presence of children is associated with a higher propensity to refinance at the end of the fixed deal across all specifications. The average marginal effect (AME) reported in the table is positive though not significant in the specification without controls. While the sample size is limited, these results show for a period when differences in reference points across fixation periods are negligible on average, borrowers that are expected to have a preference for longer fixation periods do not show a tendency to refinance less, rather they appear to be relatively more active. This evidence mitigates the concern that unobserved borrower characteristics associated with the choice of the fixation period are confounding my results.

[14] The value of the house is the one reported to the question: "About how much would you expect to get from your main home if you sold it today?". The question is asked only if the respondent is mainly or jointly responsible for financial decision making.

Table 1.7: Households with Children

This table reports the average marginal effect (AME) estimates from a logit regression of a dummy variable indicating whether the mortgage has been refinanced after the end of the fixed deal on a dummy for whether there are children (aged less than 15 years old) living in the household. Education is a discrete variable that measures the level of respondent's formal educational attainment. tstatistics are reported in parentheses. Key: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

	Average Marginal Effect				
Dependent Variable: Refinance Dummy	(1)	(2)	(3)	(4)	(5)
With Children	0.069	0.112*	0.145**	0.117*	0.270**
	(1.485)	(1.822)	(2.214)	(1.653)	(3.268)
Controls					
Age		✓	✓	✓	✓
Loan-to-Value		✓	✓	✓	✓
Total Debt-to-Income			✓	✓	✓
Log Income				✓	✓
Education					✓
Observations	218	160	140	115	77

The Effect of Maturity when Reference Points are similar

In this section, I propose a way to use the loan-level data to test whether differences in fixation period are related to lower refinancing activity for reasons other than through the effect on reference points. I show that when borrowers have similar reference rates at reset, a longer fixation period is not associated with a lower probability of observing refinancing.

Even though average interest rates on 5-years contracts are consistently higher than those on 2-years contracts over the period considered in this analysis, there is still considerable variation in the rate spread between 5-years and 2-years mortgages resetting at the same time. Even though the same rate applies to all borrowers that select into the same product, term premia, i.e. the rate differential charged to fix the rate for 5 years rather than 2 years, vary across lenders, time and loan leverage. I exploit this variation to compare the refinancing rates across borrowers with different fixation periods as a function of the difference in refinancing rates. In particular, under my key identifying assumption, if rates for long and short fixation contracts are similar, I should not observe significant differences in refinancing rates.

For the subsample of loans for which I can determine the length of the expired introductory period, I compute the difference in reference rates between 5-years and 2-years mortgages within lender, reset month and LTV class. I then estimate the effect of having a longer fixation period on the probability to refinance in subsamples based on the difference in average reference rates. Estimation results are reported in Table 1.8. In the first column the effect is estimated only considers observations where the average difference in reset rates is small, namely below the 10th percentile

Table 1.8: Isolating Fixation Period Choice from Reference Point Differences

For this regression I exploit the variation in average rates on 5-years and 2-years mortgages across lender*event month*LTVbin clusters to test whether fixation periods affect refinancing decisions beyond the effect through reference points. I first compute the average expired fixed rate over mortgages resetting in the same month, with the same lender and within the same LTV range for 5-years loans and for 2-years loans separately. The difference $\bar{r}^{5yr}_{l,t,LTV} - \bar{r}^{2yr}_{l,t,LTV}$ between these averages reflects how far apart reference points of mortgages with different fixation period are in a given cluster. I run a regression of $Refinance_i$ on an indicator for whether the loan is a 5-years fixed deal on subsamples based on the distribution of this measure. In the columns headers, $p10$, $p25$, $p50$ and $p75$ are the 10th, 25th, 50th and 75th percentile, respectively. Reset events in the $< p10$ group may differ in terms of fixation period but have fairly homogeneous reference points. Regressions include controls for the log of the loan balance at reset, remaining term (in years), borrower's age, log income and dummies for whether the mortgage is a repayment mortgage, the borrower is a first-time buyer and the income was verified upon origination. All specifications also include month, lender and LTV bin fixed effects. tstatistics are reported in parentheses. Key: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

	Cutpoints of the distribution of $(\bar{r}^{5yr}_t - \bar{r}^{2yr}_t)$				
	$< p10$	$< p25$	$> p25$	$> p50$	$> p75$
LHS Variable: $Refinance_i$	(1)	(2)	(3)	(4)	(5)
5-yrs dummy	0.010	-0.040***	-0.122***	-0.141***	-0.175***
	(0.578)	(-2.746)	(-14.188)	(-12.148)	(-9.465)
Controls	✓	✓	✓	✓	✓
Month FE	✓	✓	✓	✓	✓
Lender FE	✓	✓	✓	✓	✓
LTV bin FE	✓	✓	✓	✓	✓
Observations	1,831	3,998	12,585	9,171	4,235
R-squared	0.078	0.127	0.131	0.113	0.144

of the distribution. The coefficient on the 5-years dummy shows that there is no significant difference in propensity to refinance across borrowers with different fixation periods. Columns 2-5 repeat the same exercise over different samples based on the difference in reference rates. The results show that the larger the difference in reference rates, the more negative the coefficient, i.e. the larger the difference in refinancing frequency between borrowers on a 5-years contract versus those on a 2-years contract. These results confirm our hypothesis that differences in fixation periods lead to different refinancing probabilities through the effect on the reference rate. Instead, I find no evidence that borrowers that prefer to bear less interest rate risk are also less inclined to refinance at the end of the introductory period through other channels.

1.7.2 Sample Selection Issues and Learning from Repeated Refinancing

I address concerns about learning from repeated refinancing by restricting the analysis to a sample of borrowers that have little or no experience with refinancing. When I look at borrowers that face a rate reset for the first time I simultanously take care of the sample selection issue.

The data contains a variable that indicates the purpose of the mortgage. The two most common values are *Purchase* and *Remortgage*. Remortgage indicates that the borrower previously had a mortgage with a different lender and decided at some point to remortgage with the current lender (so called external remortgagors). Everytime that the intorductory deal ends, a borrower can decide to stay with her current lender or remortgage to a new one. More frequent resets give more opportunities to shop around and take advantage of the lowest rates on the market. As a first way to control for experience with refinancing, in Column (1) of Table 1.9 I run the baseline specification excluding external remortgagors. The coefficient on the dummy GAIN ($D_i > 0$) indicates that reference dependence affects refinancing decisions also among borrowers that were always with the same lender.

I then restrict the analysis to borrowers that face a rate reset for the first time. I take two different approaches to identify mortgages that did not reset before. First, while the field indicating the purpose of the loan is supposed to be static according to the ED guidelines to the data providers, two lenders report it dynamically, changing it from purchase to remortgage when borrowers refinance internally, i.e. with the same lender, and therefore keep being reported afterwards. This lucky occurrence allows me to clearly identify first time refinancers for these two mortgage lenders. Results are reported in column (2). I then identify first time resets among mortgages that have been originated at most five years before the observed reset event and for which I observe at least two years of submissions prior to the event.

The estimated coefficient is again negative and significant, suggesting that my findings are not driven by sample selection issues or differences in how experienced borrower are with the refinancing process that correlate with the choice of the fixation period.

1.7.3 Self-selection into Maturities based on Private Information

To control for the incentive of high-risk borrowers self-selecting into longer maturities, I run a regression similar to the one in 1.5 where I additionally control for the difference in term premia between 5-years and 2-years maturities at the time of the origination of the loan.[15] Specifically, I estimate the following specification

$$Refi_i = \alpha + \beta D_i + \gamma_1 \mathbb{1}(2yrs) TermPremium_{t-2} + \gamma_2 \mathbb{1}(5yrs) TermPremium_{t-5} + \gamma_3 W_i + \varepsilon_i \tag{1.6}$$

where $\mathbb{1}(2yrs)$ in a dummy variable that equals 1 if loan i is a 2-years fixed rate deal and $TermPremium_{t-2}$ is the term premium two years prior to the reset. $\mathbb{1}(5yrs)$ and $TermPremium_{t-5}$ are defined similarly for 5-years mortgages. W is a set of controls. Mortgages are approximately classified into 2 and 5 years fixations based on whether the interest rate on the resetting mortgage is below or above the midpoint of average

[15] Appendix 1.9.2 describes the estimation procedure for term premia.

Table 1.9: Evidence from the First Reset Date

This Table shows regression results for three different subsamples of borrowers. Column (1) excludes external remortgagors, i.e. borrowers that were previously with another lenders and that refinanced to the current one before entering the dataset. Columns (2) and (3) take two different approaches to identify mortgages that reset for the first time.

	Exclude External Remortgagors	First Time Resets	
		Reported Directly	Identified from the time series
	(1)	(2)	(3)
GAIN dummy	-0.128***	-0.094***	-0.090***
	(-22.871)	(-3.948)	(-7.493)
Controls	✓	✓	✓
Lender Fe	✓	✓	✓
Reset Month FE	✓	✓	✓
LTV bin FE	✓	✓	✓
Observations	45,647	2,527	15,898
R-squared	0.138	0.081	0.123

interest rates two and five years prior. This specification makes sure to control for the term premium that is relevant at the moment of the origination of the loan.

Estimation results are reported in Table 1.10. Under the assumption that borrower do self-select based on their privately known risk and that refinancing is denied as a function of this, we would expect γ_1 positive and γ_2 negative. Everything else constant, the higher the premium of 5-years contracts over 2-years contracts at $t-2$ indicates a better risk-pool of resetting borrowers with 2-years loans. The same intuition but in the opposite direction, is true for the term premium at $t-5$. Columns (1)-(2) simply show that this may be a plausible explanation. However, when in columns (3)-(4) I include the distance from the reference point D, the coefficients on the gammas turn insignificant and are of the wrong sign. Moreover, the coefficient on D remains negative and significant, indicating that self-selection is unlikely to explain difference in refinancing behavior. In columns (5)-(6) I repeat the analysis using moving averages of term premia, computed over a rolling 6-month window to control for some of the noise in the estimation of term premia.

1.8 Conclusion

Using loan-level data on fixed rate mortgages that automatically reset to reversion rates on pre-determined dates, I present evidence that borrowers' evaluate the benefits from refinancing relative to reference points that determine whether failing to refinance is perceived as a loss or as a gain. I show that perceived changes in monthly mortgage payments significantly affect the probability of refinancing in the cross-section of borrowers, imposing a severe friction to the pass-through of monetary

Table 1.10: Addressing Loan Maturity as a Screening Device

This table presents estimation results from the regression specification in equation (1.6). The dependent variable is a 0/1 indicator for whether the mortgage is refinanced within six months from the end of the initial fixed period. The distance from the reference rate D, defined for each individual borrower facing a rate reset as $D_i = r_i^R - SVR_i$ (in percentage points). r_i^R is the expired fixed rate of mortgage i and SVR_i is the corresponding reversion rate, i.e. the SVR of loan's lender l. $\overline{r^{2yr}}$ is the average rate on a newly originated 2-years initial fixed deal at the time of the reset. $\mathbb{1}(2yrs)$ in a dummy variable that equals 1 if loan i is a 2-years fixed rate deal and $TermPremium_{t-2}$ is the term premium two years prior to the reset. $MATermPremium_{t-2}$ is a moving average computed over a rolling window of 6 months. The first set of control variables includes log balance, years until maturity, age of the borrower, log income and current LTV. The second set of controls corresponds to a repayment method dummy, an income verification dummy and a dummy for whether the borrower is a first-time buyer. Reported t-statistics in parentheses are clustered at the month when the rate resets and at the region (nuts-2) where the property is located. Key: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
D_i (relative gain)	-0.066***	-0.085***	-0.092***	-0.090***	-0.078***	-0.067***	-0.077***
	(-9.652)	(-10.975)	(-8.602)	(-9.277)	(-9.537)	(-7.428)	(-7.970)
$\mathbb{1}$(2yrs) × Term Premium$_{t-2}$		-1.479	-2.771*	-2.879*	-1.095		
		(-1.431)	(-1.954)	(-1.867)	(-1.236)		
$\mathbb{1}$(5yrs) × Term Premium$_{t-5}$		0.734	0.515	0.373	-0.124		
		(1.404)	(1.647)	(1.332)	(-0.496)		
$\mathbb{1}$(2yrs) × MA Term Premium$_{t-2}$						-3.856***	-2.334***
						(-3.546)	(-3.047)
$\mathbb{1}$(5yrs) × MA Term Premium$_{t-5}$						-2.305**	-0.805**
						(-2.453)	(-2.267)
$SVR_{it} - \overline{r_t^{2yr}}$	0.070**	0.065**	0.217**	0.211***	0.081**	0.059*	0.082**
	(2.100)	(2.077)	(2.680)	(2.869)	(2.159)	(2.004)	(2.170)
$slope_t$			12.100	12.218			
			(1.667)	(1.692)			
$bondpremium_t$				-1.960			
				(-0.694)			
Constant	-0.765***	-0.738***	-1.220***	-1.197***	-0.737***	-0.671***	-0.720***
	(-6.322)	(-6.710)	(-4.350)	(-4.737)	(-5.646)	(-6.570)	(-5.437)
Controls	✓	✓	✓	✓	✓	✓	✓
Lender FE	✓	✓	✓	✓	✓	✓	✓
LTV bin FE	✓	✓	✓	✓	✓	✓	✓
Month FE					✓		✓
Observations	66,486	66,486	66,486	66,486	66,486	66,486	66,486
R-squared	0.139	0.140	0.144	0.144	0.156	0.143	0.156

policy relative to a rational model of mortgage refinancing in which refinancing responds solely to the future effective cost of borrowing. While average refinancing rates positively correlate with proxies for financial literacy, reference point effects are robust and stable in magnitude across income and age groups, thus providing a possibile explanation for heterogeneity in refinancing choices across similar borrowers. Overall, my results show that the current design of mortgages in the UK implies that the ability of the central bank to stimulate household consumption in a recession depends on the path of interest rates in the past through the distribution of households reference points and on reversion rates set by lenders with market power on their current customers.

References

Abel, Joshua and Andreas Fuster (2018). "How do mortgage refinances affect debt, default, and spending? Evidence from HARP". In: *Evidence from HARP (February, 2018). FRB of New York Staff Report* 841.

Agarwal, Sumit, John C Driscoll, and David I Laibson (2013). "Optimal Mortgage Refinancing: A Closed-Form Solution". In: *Journal of Money, Credit and Banking* 45.4, pp. 591–622.

Agarwal, Sumit, Richard J Rosen, and Vincent Yao (2016). "Why do borrowers make mortgage refinancing mistakes?" In: *Management Science* 62.12, pp. 3494–3509.

Agarwal, Sumit et al. (2015a). "Do consumers choose the right credit contracts?" In: *The Review of Corporate Finance Studies* 4.2, pp. 239–257.

Agarwal, Sumit et al. (2015b). *Mortgage refinancing, consumer spending, and competition: Evidence from the home affordable refinancing program*. Tech. rep. National Bureau of Economic Research.

Andersen, Steffen et al. (2015). *Inattention and inertia in household finance: Evidence from the Danish mortgage market*. Tech. rep. National Bureau of Economic Research.

Andersen, Steffen et al. (2019). "Reference Points in the Housing Market". In: *Available at SSRN 3396506*.

Auclert, Adrien (2019). "Monetary policy and the redistribution channel". In: *American Economic Review* 109.6, pp. 2333–67.

Bajo, Emanuele and Massimiliano Barbi (2018). "Financial illiteracy and mortgage refinancing decisions". In: *Journal of Banking & Finance* 94, pp. 279–296.

Baker, Malcolm, Xin Pan, and Jeffrey Wurgler (2012). "The effect of reference point prices on mergers and acquisitions". In: *Journal of Financial Economics* 106.1, pp. 49–71.

Barberis, Nicholas C (2013). "Thirty years of prospect theory in economics: A review and assessment". In: *Journal of Economic Perspectives* 27.1, pp. 173–96.

Beraja, Martin et al. (2018). "Regional Heterogeneity and the Refinancing Channel of Monetary Policy". In: *The Quarterly Journal of Economics* 134.1, pp. 109–183.

Berger, David W et al. (2018). *Mortgage Prepayment and Path-Dependent Effects of Monetary Policy*. Tech. rep. National Bureau of Economic Research.

Beshears, John et al. (2009). "The importance of default options for retirement saving outcomes: Evidence from the United States". In: *Social security policy in a changing environment*. University of Chicago Press, pp. 167–195.

Beshears, John et al. (2015). "The effect of providing peer information on retirement savings decisions". In: *The Journal of finance* 70.3, pp. 1161–1201.

Beshears, John et al. (2018). *Behavioral household finance*. Tech. rep. National Bureau of Economic Research.

Best, Michael Carlos et al. (2018). *Estimating the elasticity of intertemporal substitution using mortgage notches*. Tech. rep. National Bureau of Economic Research.

Campbell, John Y (2006). "Household finance". In: *The journal of finance* 61.4, pp. 1553–1604.

Campbell, John Y and Joao F Cocco (2003). "Household risk management and optimal mortgage choice". In: *The Quarterly Journal of Economics* 118.4, pp. 1449–1494.

Cloyne, James et al. (2019). "The Effect of House Prices on Household Horrowing: A New Approach". In: *American Economic Review*.

Cochrane, John H and Monika Piazzesi (2005). "Bond risk premia". In: *American Economic Review* 95.1, pp. 138–160.

DeFusco, Anthony and John Mondragon (2018). "No job, no money, no refi: Frictions to refinancing in a recession". In:

DellaVigna, Stefano (2009). "Psychology and economics: Evidence from the field". In: *Journal of Economic literature* 47.2, pp. 315–72.

Di Maggio, Marco, Amir Kermani, and Christopher Palmer (2016). *How quantitative easing works: Evidence on the refinancing channel*. Tech. rep. National Bureau of Economic Research.

Di Maggio, Marco et al. (2017). "Interest rate pass-through: Mortgage rates, household consumption, and voluntary deleveraging". In: *American Economic Review* 107.11, pp. 3550–88.

Dougal, Casey et al. (2015). "Anchoring on credit spreads". In: *The Journal of Finance* 70.3, pp. 1039–1080.

Eichenbaum, Martin, Sergio Rebelo, and Arlene Wong (2018). "State dependent effects of monetary policy: The refinancing channel". In:

Financial Times (2017). "Don't fall victim to your lenders' standard variable rate". In: March 8.

Fuster, Andreas and James Vickery (2014). "Securitization and the fixed-rate mortgage". In: *The Review of Financial Studies* 28.1, pp. 176–211.

Fuster, Andreas and Paul S Willen (2017). "Payment size, negative equity, and mortgage default". In: *American Economic Journal: Economic Policy* 9.4, pp. 167–91.

Genesove, David and Christopher Mayer (2001). "Loss Aversion and Seller Behavior: Evidence from the Housing Market". In: *The Quarterly Journal of Economics* 116.4, pp. 1233–1260.

Goggin, Jean et al. (2012). "The financial crisis and the pricing of interest rates in the Irish mortgage market: 2003-2011". In: *Research Technical Papers* 1.

Greenwald, Daniel (2018). "The mortgage credit channel of macroeconomic transmission". In:

Hertzberg, Andrew, Andres Liberman, and Daniel Paravisini (2018). "Screening on loan terms: evidence from maturity choice in consumer credit". In: *The Review of Financial Studies* 31.9, pp. 3532–3567.

Hirshleifer, David (2015). "Behavioral finance". In: *Annual Review of Financial Economics* 7, pp. 133–159.

Johnson, Eric J, Stephan Meier, and Olivier Toubia (2019). "What's the Catch? Suspicion of Bank Motives and Sluggish Refinancing". In: *The Review of Financial Studies* 32.2, pp. 467–495.

Kahneman, Daniel, Jack L Knetsch, and Richard H Thaler (1991). "Anomalies: The endowment effect, loss aversion, and status quo bias". In: *Journal of Economic perspectives* 5.1, pp. 193–206.

Kahneman, Daniel and Amos Tversky (1979). "Prospect theory: An analysis of decision under risk". In: *Econometrica* 47.2, pp. 263–292.

Keys, Benjamin J, Devin G Pope, and Jaren C Pope (2016). "Failure to refinance". In: *Journal of Financial Economics* 122.3, pp. 482–499.

Koijen, Ralph SJ, Otto Van Hemert, and Stijn Van Nieuwerburgh (2009). "Mortgage timing". In: *Journal of Financial Economics* 93.2, pp. 292–324.

Kőszegi, Botond and Matthew Rabin (2006). "A model of reference-dependent preferences". In: *The Quarterly Journal of Economics* 121.4, pp. 1133–1165.

— (2007). "Reference-dependent risk attitudes". In: *American Economic Review* 97.4, pp. 1047–1073.

Loughran, Tim and Jay R Ritter (2002). "Why don't issuers get upset about leaving money on the table in IPOs?" In: *The Review of Financial Studies* 15.2, pp. 413–444.

Maturana, Gonzalo and Jordan Nickerson (2018). "Teachers Teaching Teachers: The Role of Workplace Peer Effects in Financial Decisions". In: *The Review of Financial Studies* 00.0.

Paiella, Monica and Alberto Franco Pozzolo (2007). "Choosing between fixed-and adjustable-rate mortgages". In: *Household Credit Usage*. Springer, pp. 219–236.

Pope, Devin G and Maurice E Schweitzer (2011). "Is Tiger Woods loss averse? Persistent bias in the face of experience, competition, and high stakes". In: *American Economic Review* 101.1, pp. 129–57.

Scharfstein, David and Adi Sunderam (2016). "Market power in mortgage lending and the transmission of monetary policy". In: *Unpublished working paper. Harvard University*.

The Guardian (2008). "Big banks finally back down to pass on rate cut". In: November 7.

The Telegraph (2019). "What happens when my fixed-rate mortgage ends?" In: March 6.

Tracy, Joseph and Joshua Wright (2016). "Payment changes and default risk: The impact of refinancing on expected credit losses". In: *Journal of urban Economics* 93, pp. 60–70.

Tversky, Amos and Daniel Kahneman (1981). "The framing of decisions and the psychology of choice". In: *Science* 211.4481, pp. 453–458.

Wong, Arlene (2019). "Refinancing and The Transmission of Monetary Policy to Consumption". In:

1.9 Appendix

1.9.1 Additional Material

Figure 1.7: Sample Representativeness - Borrowers Characteristics

This figure compares borrowers characteristics in the BoE/NMG Survey of Households Finances and in the ED loan-level database. I plot estimated kernel densities of household's pre-tax income on the left and of borrower's age on the right. Income is reported as of loan origination in the ED database, while it is current income in the survey. Age in the ED data is the age of the primary borrower when the fixed rate deal ends. Among survey respondents, I consider only homeowners that have a mortgage on their house. This restricts the sample to 8,880 observations with non-missing information about income and age over the period 2014-2018. The ED sample refers only to loans used in the analysis, namely hybrid loans that reset in the sample period.

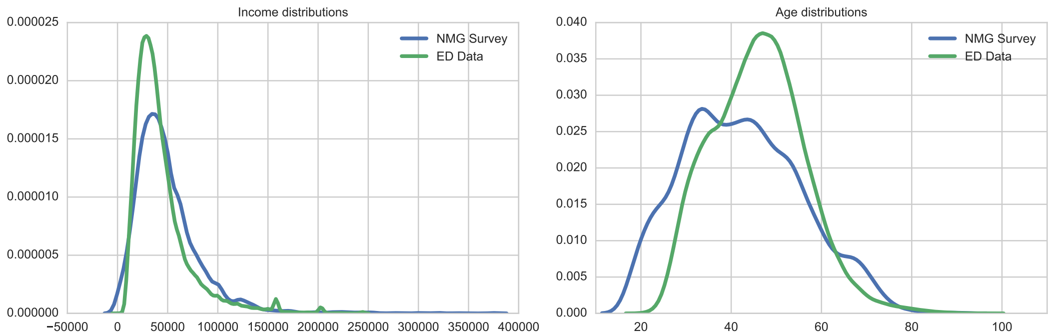

Figure 1.8: Reset Events and Rifinancing Rate

The figure provides descriptive statistics about the number of initial deal ends that we observe in our sample and of the corresponding refinancing frequency. The blue bars indicate the number of loan resets aggregated by quarter (left axis). The red line plots the average refinancing rate in each quarter (right axis).

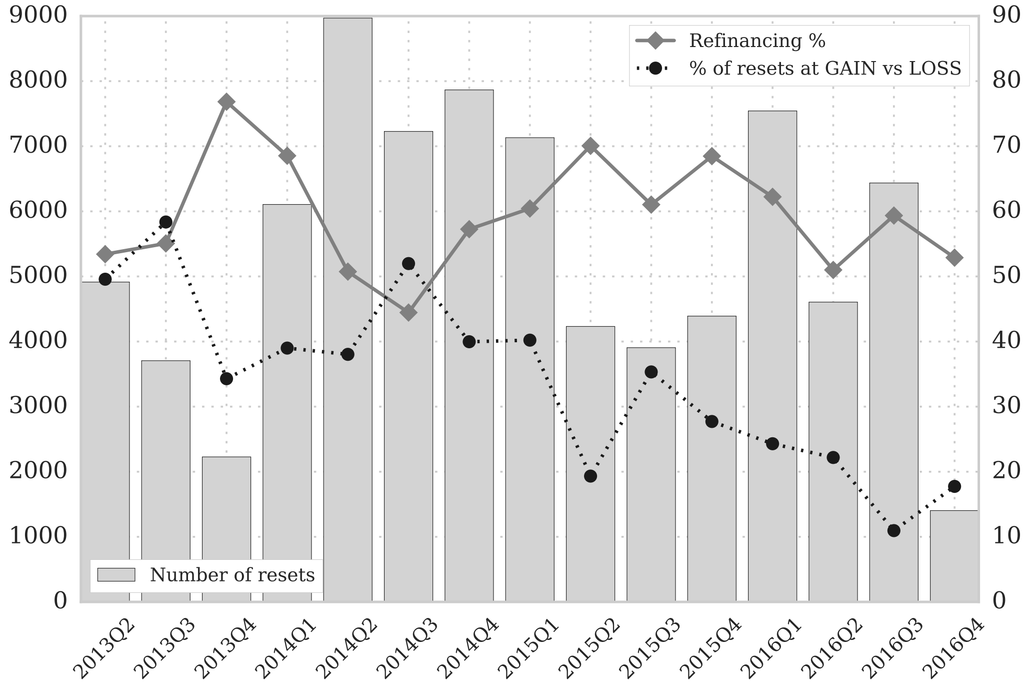

Figure 1.9: Lenders' SVRs

The figure shows UK lenders' historical standard variables rates (SVRs). This is the rate to which most fixed rate mortgages revert at the end of the introductory period.

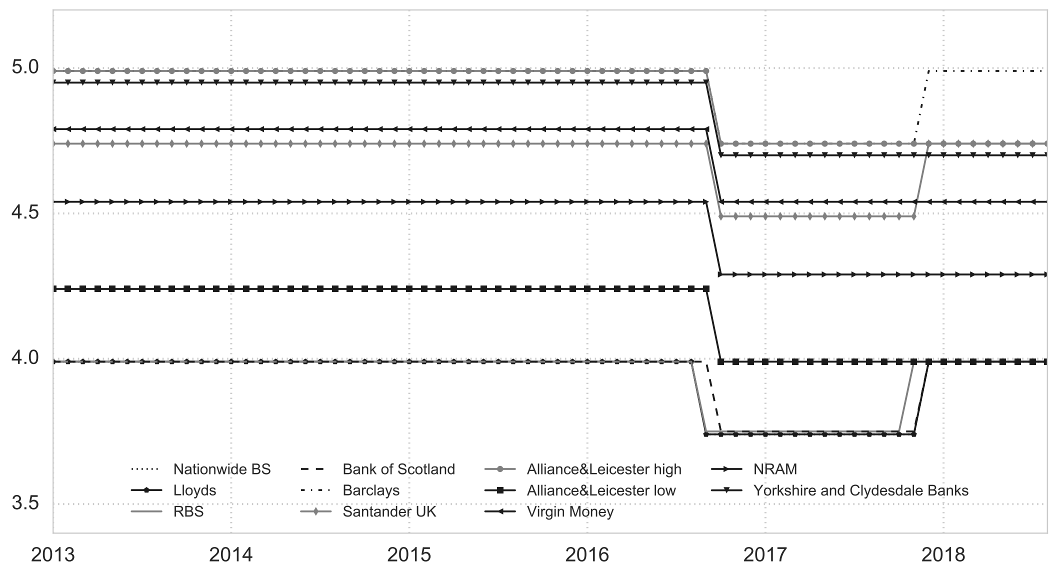

Table 1.11: Robustness for Table 1.3: Alternative Refinancing Horizons

The table shows OLS regressions like those in column 3 of Table 1.3 using different horizons over which refinancing is defined. In the first specification, the dependent variable is a 0/1 indicator for whether the mortgage is refinanced within three months from the end of the initial fixed period. The second columns reports the baseline specification where refinancing is defined over a 6-months period. The third and fourth columns look at refinancing over a period of 9 and 12 months, respectively. The main explanatory variable of interest is the distance from the reference rate D, defined for each individual borrower facing a rate reset as $D_i = r_i^R - SVR_i$ (in percentage points) and standardized. r_i^R is the expired fixed rate of mortgage i and SVR_i is the corresponding reversion rate, i.e. the SVR of loan's lender l. Reported t-statistics in parentheses are clustered at the month when the rate resets and at the region (nuts-2) where the property is located. Key: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

	(1) 3 months	(2) 6 months	(3) 9 months	(4) 12 months
D (standardized)	-0.072***	-0.078***	-0.081***	-0.084***
	(-9.113)	(-9.736)	(-8.967)	(-9.688)
Log Balance	0.092***	0.100***	0.105***	0.109***
	(9.614)	(9.567)	(10.003)	(10.897)
Years to Maturity	0.004***	0.003***	0.002**	0.002**
	(4.778)	(3.384)	(2.740)	(2.435)
Borrower's Age	-0.002***	-0.003***	-0.004***	-0.005***
	(-3.082)	(-4.048)	(-4.706)	(-5.834)
Log Income	-0.029***	-0.023***	-0.017***	-0.011*
	(-5.159)	(-4.173)	(-3.583)	(-1.988)
LTV	0.001*	0.001**	0.001*	0.001
	(1.935)	(2.175)	(2.029)	(1.186)
month x lender x ltv	Yes	Yes	Yes	Yes
Observations	84,317	79,468	71,873	58,047
R-squared	0.146	0.151	0.158	0.164
Refi Fraction	53%	59%	61%	62%

Table 1.12: Robustness for Table 1.3: Increased Sample

The table shows OLS regressions like those in columns 1-4 of Table 1.3 using a larger sample. Specifically, I add back to the sample those observations with missing borrower's age, property location, repayment type and income verification information. The right panel reports the results on the restricted sample used for Table 1.3 without including controls for borrower's age to allow comparability. The dependent variable is a 0/1 indicator for whether the mortgage is refinanced within three months from the end of the initial fixed period. The main explanatory variable of interest is the distance from the reference rate D, defined for each individual borrower facing a rate reset as $D_i = r_i^R - SVR_i$ (in percentage points) and standardized. r_i^R is the expired fixed rate of mortgage i and SVR_i is the corresponding reversion rate, i.e. the SVR of loan's lender l. Reported is the effect of one standard deviation change in D, and t-statistics in parentheses are clustered at the month when the rate resets. Key: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

	Unrestricted sample				Restricted sample			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
D (standardized)	-0.083***	-0.086***	-0.071***	-0.068***	-0.083***	-0.099***	-0.079***	-0.076***
	(-6.314)	(-13.878)	(-10.706)	(-9.320)	(-5.649)	(-13.592)	(-9.888)	(-9.385)
Log Balance			0.064***	0.064***			0.091***	0.089***
			(11.597)	(11.688)			(13.016)	(12.547)
Years to Maturity			0.006***	0.006***			0.006***	0.006***
			(10.327)	(10.095)			(6.917)	(7.147)
Log Income			-0.000	-0.003			-0.020***	-0.021***
			(-0.155)	(-0.986)			(-2.890)	(-2.994)
LTV			0.001***	0.001***			0.002***	0.002***
			(4.188)	(3.775)			(3.561)	(3.132)
month x lender x ltv	No	Yes	Yes	Yes	No	Yes	Yes	Yes
loy FE	No	No	No	Yes	No	No	No	Yes
Observations	118,608	118,608	118,608	118,608	79,467	79,467	79,467	79,467
R-squared	0.028	0.194	0.210	0.212	0.028	0.126	0.149	0.150

Table 1.13: Robustness for Table 1.3: Logit Regression

The table reports the average marginal effect (AME) estimates from a logit regression of a 0/1 indicator for whether the mortgage is refinanced within six months from the end of the initial fixed period. The main explanatory variable of interest is the distance from the reference rate D, defined for each individual borrower facing a rate reset as $D_i = r_i^R - SVR_i$ (in percentage points). r_i^R is the expired fixed rate of mortgage i and SVR_i is the corresponding reversion rate, i.e. the SVR of loan's lender l. Reported t-statistics in parentheses are clustered at the month when the rate resets. Key: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

	(1)	(2)	(3)	(4)
D	-0.065***	-0.070***	-0.065***	-0.079***
	(-5.746)	(-6.707)	(-6.353)	(-13.202)
Log Balance		0.148***	0.206***	0.171***
		(11.193)	(22.235)	(29.287)
Years to Maturity		0.002	-0.002*	-0.002***
		(1.226)	(-1.946)	(-2.751)
LTV		-0.001**	-0.002***	-0.001***
		(-2.204)	(-3.941)	(-2.734)
Borrower's Age		-0.004***	-0.003***	-0.003***
		(-5.868)	(-5.170)	(-4.985)
Log Income		-0.059***	-0.085***	-0.069***
		(-5.298)	(-8.418)	(-7.765)
Repayment Mortgage			0.168***	0.162***
			(11.560)	(12.739)
Income Verified			0.084***	-0.029***
			(3.017)	(-4.741)
First Time Buyer			-0.035**	-0.035**
			(-2.169)	(-2.437)
Month Dummies	No	No	No	Yes
Lender Dummies	No	No	No	Yes
Observations	79,468	79,468	79,468	79,463

1.9.2 Term Premia Estimation

To control for the incentives of selecting into different loan maturities based on private information about one's ability to repay (Hertzberg, Liberman, and Paravisini, 2018), I include estimated term premia at the time of the choice about the fixation period in regression (1.6). For the purpose of this exercise, I estimate risk premia from a OLS regression of excess one-year log returns of two and five maturities bonds on the level, slope and "CP factor" (Cochrane and Piazzesi, 2005). The level and slope factors are, respectively, the first and second principal component from the correlation matrix of bond yields with maturities from 1 to 10 years. As the "CP Factor" I use the predicted value from a regression of average future excess returns of two to ten-year bonds on average forward rates. Specifically, I first estimate the following regression

$$\overline{rx}_{t+1} = \gamma' \overline{f}_t + \varepsilon_{t+1} \tag{1.7}$$

I then run a regression on contemporaneous excess returns on level, slope and the predicted value from equation (1.7):

$$rx_t^{(n)} = c_1 \text{level}_t + c_2 \text{slope}_t + c_3(\widehat{\gamma}' \overline{fs}_t) + \eta_{t+1} \tag{1.8}$$

I calculate the difference in term premia of a five-year and two-year bonds as $\widehat{rx}_t^{(5)} - \widehat{rx}_t^{(2)}$. Data are from Thomson Reuters Datastream.

2 Mortgage Defaults and Positive Equity: Lessons from Europe

2.1 Introduction

Why do households default on their mortgage debt? This question has mainly been investigated by looking at borrower's behaviour in the United States, especially since the recent subprime crisis. At the same time, it is quite common to read that the European mortgage market is less risky, and has performed better than its counterpart in the U.S. either because of (1) the lack of moral hazard on the part of lenders due to a less widespread use of OTD models of origination, (2) the presence of recoursability in most European markets, and (3) lower loan-to-value ratios. However, there are many institutional differences between mortgage markets in Europe and the United States, including differences in mortgage characteristics and financing structures, which are likely to affect borrowers' behavior and loan performance and shall be controlled for.[1]

In this paper, we aim to fill this gap by empirically investigating mortgage default behavior in Europe (Euro Area plus UK) using a novel database of loan-level data that is provided by the European DataWarehouse (ED henceforth). ED is a centralised European platform for ABS loan level data. In particular, ED collects information on securitized residential mortgages backing RMBS that banks pledge as collateral in ABS transactions with the European Central Bank. The main focus is on understanding what triggers the default decision in a setting where mortgages are recourse loans (no debt forgiveness), i.e. when borrowers are responsible upon default for the difference between the value of the outstanding debt and the value of the house.

Using the loan balance as detailed in the ED data set, we identify loans with 90+ days of delinquency, which we classify as defaults. Our first finding is that a large fraction of defaults in Europe happens at positive equity: the majority of borrowers who default on their loan do so when the value of their collateral would be in principle enough to repay the debt. The estimated equity at default averages about 25% and is as large as 2.5 times an household's annual income. This evidence may seem puzzling as, from a theoretical viewpoint, a rational borrower should always prefer

[1] According to Campbell (2012) among others, the United States has much to learn from practices in Europe.

to sell the house and prepay rather than default, if equity is positive. Alternatively, the borrower may try to renegotiate the terms of the loan or to refinance it. The result that many borrowers find themselves forced to default on their loans shows that either of these options appear out of reach for a significant fraction of the European market.

The finding that European borrowers are mostly (and largely) above-water at default cast a net with much of the existing literature that has put the emphasis on understanding the extent of strategic, or "ruthless" default under non-recoursability (Bhutta, Dokko, and Shan (2017)). The standard frictionless model of Kau, Keenan, and Kim (1994) predicts that borrowers should find it optimal to stop paying and exercise their option to exit the mortgage when equity is sufficiently negative, so that the instantaneous relief from mortgage payments exceeds the option value of waiting a recovery in house prices. When taking into account the costs of default and borrowing constraints, Campbell (2012) shows that the optimal trigger level of home equity should be higher (i.e. less negative) for financially constrained households, for which the marginal utility of immediate financial relief is highest.

How does recoursability change default behavior? A common argument that is put forward is that full recourse has beneficial impact on both borrowers' quality (ex-ante), as it prevents wimpy households from entering the market, and behavior (ex-post), as it deters "strategic" default (i.e. unwillingness to pay). If the threat of lender recourse is successful, one should therefore only observe defaults that are triggered by inability to pay, i.e. only borrowers that are hit income or liquidity shocks should enter default. In this setting, it is unclear how the current value of equity should affect borrowers' decision to default. Indeed, Ghent and Kudlyak, 2011 find a reduced sensitivity of default to negative equity in recourse states.

Even without the embedded option of non-recourse mortgage contracts, however, there still exists a threshold level of equity that would trigger default. While the probability of equity becoming negative increases the value of waiting under non-recourse, postponing default is risky under recourse. As the value of the collateral decreases, the recoursability threat becomes stronger and delaying default riskier. To avoid having their private assets seized in case of foreclosure, borrowers in financial distress might optimally anticipate their decision to default in the positive equity region. If selling the house takes time, above water borrowers that are experiencing financial difficulties might decide to default now to avoid the risk of being forced into default with negative equity in the future.

This argument implies that we should observe a positive relationship between the extent of borrowing constraints and the amount of equity left on the table upon default. We test this hypothesis in the data using household's income at mortgage origination to proxy the extent of borrowing contraints and higher marginal utility of consumption today, as in Campbell (2012).

We find that equity at default is significantly negatively related with the household's income at origination. This result persists even after controlling for the loan-to-value and loan-to-income ratios, and when scaling equity by price volatility. A simple back of the envelope calculation shows that borrowers in the bottom quartile of the income distribution leave about 9% of equity on the table more than those in the top quartile, which in monetary terms corresponds to about 13'000 EUR, or as much as three-fourths of their annual income.

We address the natural concern that our estimates may be biased due the way we compute equity. First, we find that the market price we use in the computation have nearly no measurement error when compared to property valuations occasionally reported by loan originators for about one-third of the sample. Additionally, we construct equity using foreclosure prices, which can be regarded as a lower bound to the price at which households would have been able to liquidate their house upon default. Yet, we continue to observe a significant fraction of defaults at positive equity and a negative equity-income relation.

We also investigate the importance of credit rationing at the originator's level. That is, we ask if low-income borrowers are more limited in their ability to refinance or renegotiate the terms of the loan. A tightening of the lender's financial constraints, as measured by the accumulated losses in its loan portfolio, indeed makes the relation between income and equity at default more steep. This finding highlights that negative shocks in the financial slack of the financial intermediary result in low income borrowers having fewer access to measures of financial relief and being forced into default. However, equity at default is still significantly decreasing with income for loans whose originators are least constrained, which suggests that a credit channel cannot fully explain our results.

The analysis of loan defaults shows that the default threshold for low-income borrowers is higher than that of high-income borrowers facing the same loan terms. On the full sample of loans, we expect this finding to translate into a higher default sensitivity of low-income borrowers to changes in the value of collateral. We estimate a linear and probit model for the probability of default on the 42 million loan-year observations. Even after controlling for leverage and the level of equity in the prior month, we find that negative shocks to house prices have a larger impact on the default probability of low-income borrowers. Hence, shocks in the value of housing (and hence, equity) have differential default impact across the income distribution.

The remaining of the paper proceeds as follows. Section 2.2 reviews the literature on mortgage default decision. Section 2.3 describes the data and provides summary statistics. Section 2.4 presents the analysis of equity at default. The analysis of default rates is collected in Section 2.5. In Section 2.6, we test for the supply-side channel. Finally, Section 2.7 offers concluding remarks.

2.2 Literature Review

Since the turmoil in the US housing market, a large literature has emerged that attempts to understand what fuelled the credit expansion and lead to the collapse of the subprime mortgage market in February 2007. Mian and Sufi (2009) emphasize the unsustainable increase in lending to poor credit quality borrowers, suggesting a relaxation of credit standards and moral-hazard issues on the part of originators. Adelino, Schoar, and Severino (2016) challenge this standard view on the causes of the crisis and shows that the contribution to the increase in debt was shared across the entire distribution of borrowers, thus emphasizing the role of middle and high-income borrowers in pushing up default rates.

Understanding the drivers of default is important for efficient market regulation. In particular, research on default behaviour can help lenders and policy makers to decide what is best between principal reduction or temporary payment moratorium to prevent borrowers from losing their homes during crisis periods (e.g. Elul et al. (2010), Geanakoplos (2014)).

Default is usually described in the finance literature as the exercise of a real option and should therefore be observed when doing so increases their lifetime wealth. In absence of frictions and costs of default, borrowers should default on their loans as soon as equity turns negative, even if they could afford to pay (see Campbell and Cocco (2015) for a comprehensive model of mortgage default).

The literature on mortgage default has mostly focused on understanding how widespread strategic default is in practice and thus on assessing the relative importance of negative equity versus illiquidity as trigger of default. Several papers have found that ruthless default is uncommon at moderate levels of negative equity and that borrowers tend to default only when negative equity is combined with a negative income shock (Bajari, Chu, and Park, 2008; Elul et al., 2010; Foote, Gerardi, and Willen, 2008). Bhutta, Shan, and Dokko (2010) look at the defaults by subprime borrowers in the US and find that the median borrower does not default until equity falls below -62 percent of the current value of the property, suggesting that borrowers face very high costs upon default. In their sample, only 20 percent of the defaults seem to be purely driven by negative equity.

Recourse and non-recourse mortgages provide different incentives to default, but only few papers look at the effect of lender recourse on default behaviour. Ghent and Kudlyak (2011) compare default in recourse and non-recourse states in the US and find that the threat of recourse reduces the probability to default at any given level of negative equity. However, they find that unconditionally there is no difference in default rates, which runs counter the claim that lender recourse is the reason Europe has much lower default rates (e.g. Feldstein (2008)). Gete and Zecchetto (2018) show within a theoretical model that by discouraging default recourse systems magnify the impact of nominal rigidities and cause deeper and more persistent recessions.

Few papers use European mortgage data. Van Bekkum, Gabarro, and Irani (2017) use data on mortgage origination in the Netherlands to show that the relaxations of collateral eligibility criteria by the ECB in 2012 lead to an increase of lending to risky borrowers and thus a performance deterioration of loans. Acharya et al. (2017) examine how the February 2015 introduction of loan-to-value and loan-to-income limits on the issuance of residential mortgages in Ireland affected bank risk taking and household availability of credit. They show that banks most affected by the policy originated safer mortgages while they increased risk taking in corporate lending and security portfolios. Flodén et al. (2017) examine registry-based data on Swedish households and find that households that are highly indebted and have adjustable-rate mortgages respond stronger to changes in the monetary policy rate (cash-flow channel of monetary policy).

2.3 Data and variable construction

2.3.1 Loan data and sample selection

Our study makes use of loan-level data that are provided by the European DataWarehouse (ED henceforth). ED is a centralised European platform for ABS loan-level data. In particular, ED collects information on securitized residential mortgages backing RMBS that banks pledge as collateral in ABS transactions with the European Central Bank.[2] In order for an ABS to be eligible as collateral in Eurosystem refinancing operations, banks are required since January 2013 to provide detailed loan-level information regarding the pool of cash-flow generating assets at least at quarterly frequency. For each loan in the pool, banks are required to report loan, borrower and collateral characteristics at origination, as well as updated information on loan performance. For loans that defaulted or prepayed before 2013, the database reports retroactively information on the default or prepayment date, as well as the outstanding balance at default and possible accumulated recoveries[3].

We focus our analysis on the nine European countries with the largest number of loans in the ED database, namely Belgium, France, Germany, Ireland, Italy, the Netherlands, Portugal, Spain and the United Kingdom. According to the Association for Financial Markets in Europe (AFME), these are also the countries with the largest RMBS markets by outstanding volume in Europe and account for 53% of the

[2]See Van Bekkum, Gabarro, and Irani (2017) and Ertan, Loumioti, and Wittenberg-Moerman (2017) for other studies that use ED data.

[3]To be precise, this information is mandatory for all loans that are active in the pool at the time of submission. Banks are also strongly encouraged to submit information for all inactive loans, i.e. loans for which no cash-flows are expected in the future. In the database there are more than 3.4 million loans that are reported as inactive at first submission. For the average (median) inactive loan 47% (58%) of the mandatory fields are populated. Compared with the average (median) reporting ratio of 78% (82%) for active loans, this suggests that banks are quite keen to provide loan-level information for all loans in a pool.

Figure 2.1: Coverage of ED Data.

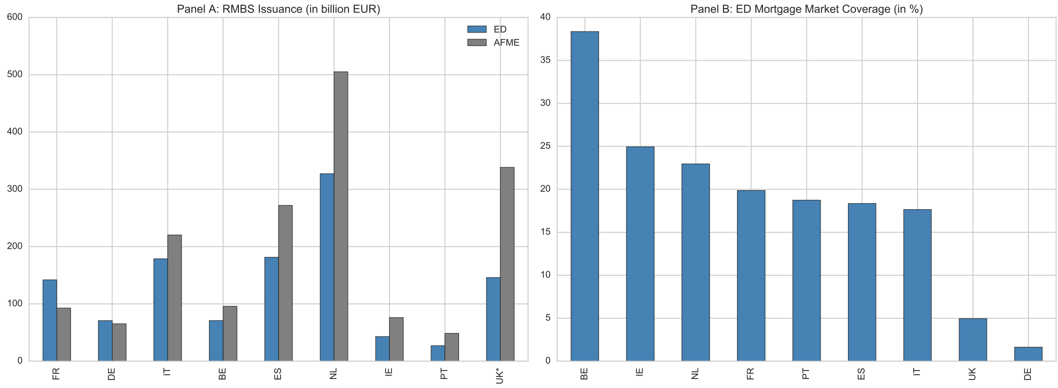

This figure provides an overview of the share of the RMBS market (Panel A) and of the underlying mortgage market (Panel B) covered by the ED database. Next to loan-level characteristics and performance updates, originators must submit to ED also information at tranche and security level for all (eligible and non-eligible) tranches of the RMBS deals. In Panel A, we plot the aggregate RMBS issuance volume by country against the data from the Association for Financial Markets in Europe (AFME). Since AFME started reporting RMBS issuance volumes in 2005, we consider RMBS issued between 2005 and 2015. For the UK we only consider the period 2009-2015 since pre-2009 RMBS vintages are almost never reported. Panel B plots the aggregate mortgage origination volume in ED as a percentage of the volume reported by the European Mortgage Federation (EMF). We only consider mortgages originated between 2004 and 2015, which corresponds to the period covered by EMF data.

overall European market. The database contains more than 12 million residential mortgages originated from January 2000 until December 2015,[4].

In Figure 2.1, we collect statistics on the coverage of ED data for the RMBS and mortgage market. In Panel A, we compare the aggregate issuance volume at country level with the one reported by the Association for Financial Markets in Europe (AFME) for the period 2005-2015. We can compute the total RMBS issuance volume in ED because originators are required to provide information on issue date and issued amount of all the tranches (eligible and non-eligible) of a submitted RMBS deal. As we can see, ED data reflect on average about 78% of a country's RMBS total issuance volume (AFME).

The mortgages reported to ED also represent a relevant fraction of the overall issued mortgages in most European countries. For the period 2004-2005, we plot in Panel B the aggregate mortgage origination volume in ED as a percentage of the volume reported by the European Mortgage Federation (EMF). This fraction is around 20% for most countries, and grows to in Belgium. This figure drops below 5% in the UK and Germany. For Germany, this figure is consistent with the evidence that most German mortgages are financed through covered bonds (*Pfandbriefe*) rather than RMBS.[5] For completeness, we keep Germany in the analysis but results are unchanged if we exclude it from the sample.

[4]To exclude years with few originations from our sample, we do not consider loans originated before 2000. This choice ensures that we can compute the regional semiannual HPI index based on the property valuation at origination contained in our data for most regions.

[5]See the European Corporate Bond Fact Book 2016 by the European Covered Bond Council available at https://hypo.org/ecbc/publications/.

For each loan in the sample, ED provides an identifier for the mortgage originator. This might be either the name of the originator of the loan or an anonymous alphanumeric code. Our sample contains the main mortgage providers of each country as well as smaller ones. We compute the aggregate volume of originated mortgages in the sample period by originator id. We note that even though some providers account for a large fraction of the loans, originations are not concentrated in few providers.

We apply a series of filters to the raw data set. Specifically, we require information on loan amount, loan maturity, property valuation at origination, household income as well as origination date to be non-missing. Moreover, we filter out loans with missing detailed geographical information on the property location.[6] Within each country, we trim household income, property valuation and loan amount to the 1st and 99th percentile. We also trim LTV and LTI ratios to get rid of reporting errors. Finally, we restrict our attention to purchase loans[7] and eliminate loans that get repurchased by the seller at any point during their lifetime.

Borrowers can take on multiple small size loans with possibly different characteristics (maturity, interest rates, repayment method, etc.) to finance the purchase of a property. This practice is common in the Netherlands and is sometimes observed in other countries as well. The observation unit in ED is a loan part. In order not to overestimate the equity of a borrower into a property it is important that we conduct our analysis at mortgage level. As in Van Bekkum, Gabarro, and Irani, 2017, we therefore aggregate loans originated at the same time by a single borrower on a single property into a single mortgage. The resulting mortgage size is the sum of the principal of its loan parts, while loan characteristics are computed as mortgage balance-weighted averages.

Table 2.1 collects summary statistics for the resulting sample. Panel A reports the total loan amount and number. We rank countries based on the total dollar volume of loans issued. The largest volume is observed for France, totaling about 190 billion EUR from about 1.6 million loans. The smallest markets are Portugal and especially Germany. For the sake of completeness we keep German data in the sample, but our main results remain unaffected if we instead exclude it. The overall amount issued of more than 760 billion EUR captures a wealth of outstanding European households' mortgage debt.

In Panel B of the table, we look at loan characteristics at origination averaged across all loans. The average loan amount issued is about 137k EUR, against more than

[6]Banks are required to report information on the geographic location of the property backing each loan. There is some heterogeneity in the reporting practices, whereas some banks report the zip code (at least 2 digits), while other report a nuts code or the name of the city. In order to obtain comparable geographical areas, we map the available information to nuts 3 codes. Nuts 3 regions therefore constitute the smallest geographical unit in our dataset. The location is identified for more than 90% of the loans in our sample in every country except for UK (56%).

[7]The purpose of the loan is a mandatory field and includes purchase, construction, renovation, equity release, debt consolidation and remortgage. For all countries, purchase loans are the most frequent and account for about two-thirds of the total.

199k EUR for the value of collateral. The resulting average loan-to-value (LTV) ratio is about 73%. However, we note significant dispersion among countries, with values of LTV ranging from 64% (Italy) to more than 85% for Germany and the Netherlands. Even more pronounced differences are observed for the ratio of loan to annual income (LTI), which averages about 4.1 but varies from as low as 2.8 (for Germany) to as high as 6.3 (for Portugal). The average loan maturity is about 24 years, with loans from Portugal being the longest at 34 years.

Overall, about 53% of the loans are fixed-rate. There is, however, very limited within-country variation in interest rate types. Four countries exhibit nearly all fixed-rate loans (these are France, the Netherlands, Belgium, and Germany; henceforth, FRM countries) while in other four countries loans are almost exclusively floating (namely, Spain, the U.K., Ireland, and Portugal; henceforth, ARM countries). Only in the case of Italy do we observe a significant fraction of either type, with about 23% loans being fixed-rate.

An additional loan characteristics, the number of lien is optional in ED's reporting template. The vast majority of loans with non-missing lien (58%) are first-lien.

Finally, Panel C collects statistics for borrower's characteristics. Average income is lowest in Portugal, Italy, and Spain, while it exceeds 60k EUR in Ireland and the U.K. The average borrower is 38 years old, and is employed for about 70% of the loans. About half of the loans contain information on the number of borrowers that are responsible for the mortgage. This number averages at 1.6, with about 43% of such loans having a single borrower.[8]

2.3.2 Variables definition

Our focus is on the relationship between a borrower's decision to default on a mortgage and the level of equity. To this end, we define a loan to be delinquent at a certain date t if payments are reported 90 or more days late for two consecutive quarters[9], or if the loan is reported as in foreclosure or in default. We then define the "time of default" as the month when the loan reaches the 90+ day delinquency mark for the first time or as the date when it is first reported in default.

Figure 2.2 displays time series of the aggregate loan originations in bn Euro (Panel a), and the resulting average default rates (Panel b). Originations increase exponentially peaking at about 100 bn Euro in 2006, and then experience a marked contraction, with the exception of a rebound in 2011-2012. Non surprisingly, a similar pattern is observed for housing prices (in red for the OECD series). When applying our definition of default, we obtain aggregate default rates in the 0 to 1% range. Unlike

[8] In case of multiple borrowers, banks can report the primary and the secondary income as two separate fields. We refer to the sum of these two income values as household income. LTI at origination is computed as the ratio of the face value of the loan to the income at household level.

[9] Two consecutive submissions might be one month or one quarter apart, depending on the submission frequency of a specific deal. We consider two consecutive quarters in order for our definition of default not to depend on the submission practice.

Table 2.1: Summary statistics

This table presents summary statistics for the loan-level dataset from ED for each country (in decreasing order of loan amount issued) and for the pooled panel (last column). Panel A reports the total loan amount and number. Panel B reports averages of loan characteristics. LTV is loan-to-value, LTI is loan-to-income, FRM is a dummy that equals one if the loan is fixed rate. Panel C collects averages of borrower characteristics. ED data contain mortgage originator identifiers, which might either be the name of the originator of the loan or an anonymous alpha-numeric code. In panel D we report the market share of the three largest originators by aggregate originated volume in each country as well as the aggregate market share of the five largest originators. An asterisk denotes fields that are optional in the reporting templates to ED and are therefore missing for a subsample of loans. Sample period is 2000 to 2015.

	FR	ES	NL	IT	BE	UK	IE	PT	DE	*Total*
Panel A: Sample characteristics										
Issued loan amount (bn EUR)	189.8	153.8	149.8	102.5	59.8	42.5	31.5	22.8	9.2	*761.7*
Number of loans (thousands)	1’577	1’047	773	888	475	348	149	229	83	*5’570*
Panel B: Loan characteristics										
LTV (%)	71.1	73.7	85.1	63.6	75.3	70.8	75.4	77.6	87.6	*73.3*
LTI	3.1	5.9	4.0	4.3	3.3	3.0	4.0	6.3	2.8	*4.1*
Loan size (EUR)	120’368	146’847	193’711	115’359	125’880	121’898	211’463	99’458	112’042	*136’751*
Collateral value (EUR)	195’475	206’954	236’797	191’717	180’738	180’904	289’824	133’937	133’234	*199’674*
Loan maturity (years)	18	28	30	22	20	23	28	34	24	*24*
Interest rate (%)	3.1	2.8	4.2	2.0	2.8	3.4	2.7	1.0	3.1	*2.9*
FRM (%)	93.4	3.2	90.5	23.2	95.0	0.1	3.9	1.4	98.7	*53.0*
Interest only (%)	0.1	0.0	42.9	0.0	0.1	17.4	5.8	0.4	4.3	*7.3*
Mortgage loan parts	1.3	1.0	1.9	1.0	1.1	1.0	1.0	1.1	1.1	*1.2*
Lien *	1.4	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	*1.1*
Panel C: Borrower characteristics										
Household’s income (EUR)	45’511	32’609	51’588	31’255	47’649	59’147	58’766	22’226	52’242	*42’187*
Age	38	36	40	39	36	38	34	34	37	*38*
Employed (%)	66.3	66.7	80.3	66.8	70.8	53.9	81.6	66.5	84.7	*68.7*
Selfemployed (%)	7.4	9.9	5.4	15.7	11.2	6.9	10.4	9.0	8.5	*9.4*
Civil/government servant (%)	19.4	9.3	0.0	2.0	11.4	1.4	0.0	14.5	5.4	*9.3*
Pensioner (%)	3.4	2.0	4.3	5.2	1.0	0.5	0.1	2.0	0.6	*3.0*
Resident (%)	100.0	61.0	100.0	100.0	100.0	99.9		97.4	100.0	*86.8*
Number of borrowers *	1.5	1.6	1.5	1.6	1.6	1.6	1.7	1.6	1.6	*1.6*
Panel D: Originators market share (%)										
First largest	12.78	15.31	12.75	19.92	40.80	30.69	42.53	18.94	43.68	
Second largest	7.14	9.36	9.43	11.60	14.90	17.84	20.53	18.16	32.23	
Third largest	7.08	8.94	7.85	7.75	11.46	16.68	14.63	14.22	21.80	
Sum 1-5	33.98	48.28	43.48	50.11	87.14	79.75	99.76	73.91	100.00	

the U.S., however, for which mortgage delinquencies peaked in 2010, delinquencies in Europe hit their maximimum in 2013 during the European sovereign crisis, and then revert back to their pre-crisis level. This timing of residential defaults lines up with that of unemployment rate (dotted line in the figure), which serves as our main control for time-varying economic conditions at the regional level.

Given our definition of default, we construct our measure of equity at default for a given loan i as follows. Let V_{i,t_0} be the value of the property at the time of loan origination. Let H_{i,t_0} be the house price index at the same time for the county (nuts 3) where the property of mortgage i is located, computed as average property value across all loans in the region. We estimate the value of the property at a given time t as:

$$\widehat{V}_{i,t} = V_{i,t_0} \frac{H_{i,t}}{H_{i,t_0}} \tag{2.1}$$

We then compute the percentage equity stake of the borrower at the time of default as:

$$\widehat{E}_{i,t} = \frac{\widehat{V}_{i,t} - B_{i,t}}{\widehat{V}_{i,t}} \tag{2.2}$$

Figure 2.2: Loan Originations and Defaults

(a) Aggregate Originations

(b) Average Default Rates (%)

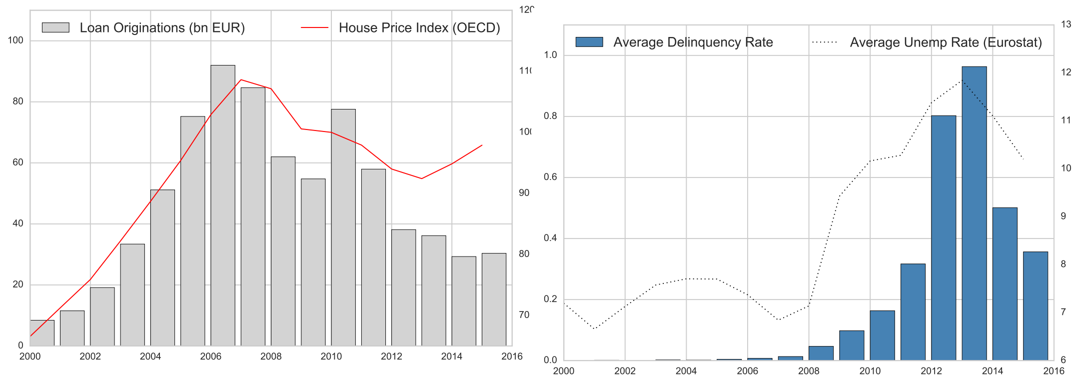

This figures shows the aggregate time series of loan originations and of default rates in our dataset. In panel (a), the grey bars indicate loan originations in billions EUR computed as the sum of the principal value of all mortgages originated in one year. The red line is the average OECD annual house price index across the countries in the sample. In panel (b), the blue bars indicate annual average default rates in percentage calculated as simple averages across countries. Default rates are computed as the ratio of loans in default based on our definition above over the total number of loans outstanding. Loans are considered outstanding from the origination until maturity, unless they become delinquent at a certain point in which case they are removed from the denominator in every subsequent period. Prepaid loans are kept in the denominator over their entire contractual life. The black dotted line is the average unemployment rate from Eurostat.

where $B_{i,t}$ is the outstanding balance on loan i at time t as reported in the ED database.

To the best of our knowledge, no HPI indices are available at the county level for Europe. Therefore, we compute $H_{i,t}$ from the ED database as the average property value across all loans originated in semester t in the same county where the property of loan i is located. The property valuation corresponds to the value of the house that lenders use to compute the LTV at origination.[10]

For the purpose of computing average house values, we consider all mortgages in ED, not only purchase mortgages, in order to maximize the number of observations. The average median number of observations for a given county in a given semester is 442 across countries. To reduce noise from counties with few observations, we require a county to have at least 30 observations (loan originations) in semester t to produce a valid house price average $H_{i,t}$. In case either H_{i,t_0} or $H_{i,t}$ cannot be used for the estimation of equity, we use the OECD annual price index. We compute implied annual country-level indices and we plot them against the index from OECD in Figure 2.3. We note that the two series move quite closely together. In fact, the correlation in returns to our HPI index and the OECD series is as high as 0.80. This evidence corroborates the reliability of our HPI estimates and mitigates concerns of sample selection biases in our data.

[10]More than 80% of the valuations come from a full inspection of the property, both internal and external for the vast majority. 13% of the valuations are done by a real estate agent and very few are done through Automated Valuation Models (AVMs) or indexing.

Figure 2.3: Implied House Price Index.

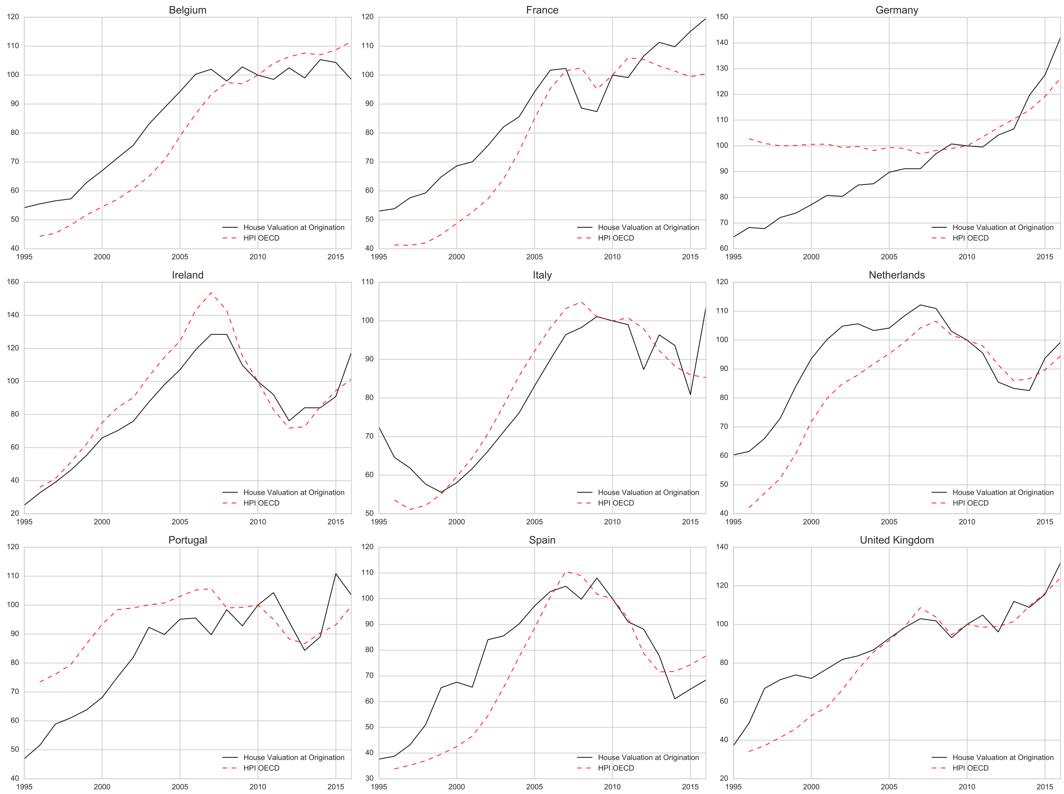

We plot the annual country-level HPI index calculated using loan originations in the ED database against the HPI index from OECD. For every year and country, the house price level is computed as the median value of the properties backing the newly originated loans. We require that currency information is non-missing, original balance non-missing and smaller than 3 million EUR, original LTV is non-missing, different from zero and smaller than 150%, For this calculation we consider all types of loan purposes.

We use $\widehat{E}_{i,t}$ as our main dependent variable in our analysis in Section 2.4. As an alternative, we also scale the amount of equity by house price volatility. This transformation allows us to verify that our analysis is not piking up correlation between house price volatility and income, which would arise if e.g. property prices were to fluctuate more in areas mostly populated by low-income borrowers. Similarly to Ghent and Kudlyak (2011), the variable is constructed as:

$$DD_{i,t} = Pr(E_{i,t} > 0) = 1 - \Phi\left(\frac{\ln B_{i,t} - \ln \widehat{V}_{i,t}}{\sigma_{i,t}}\right) \tag{2.3}$$

We compute the volatility of house prices for the county in which the loan is originated, $\sigma_{i,t}$, as the rolling standard deviation of the semi-annual HPI index return over the past 8 observations (4 years).

Table 2.2: Summary statistics for defaults

This table presents summary statistics for the sample of defaulted mortgages. Equity at default is expressed in percentage of the collateral value at the moment of default as defined in Section 2.3.3. Money left on the table is the difference between the property value and the outstanding balance at the moment of default. We compute the ratio between the money left on the table and income at origination. The remaining variables are defined as in Table 2.1.

	IT	ES	IE	PT	NL	BE	FR	UK	DE	*Total*
Panel A: Sample characteristics										
Number of defaulted mortgages	43’964	27’619	20’396	9’826	6’769	4’742	3’546	2’534	54	*119’450*
Fraction of defaults (%)	4.9	2.6	13.7	4.3	0.9	1.0	0.2	0.7	0.1	*2.1*
Equity at default (%)	49.0	21.5	1.7	33.8	-29.2	27.3	39.1	40.3	33.8	*27.4*
Percentage of defaults with positive equity	96.3	70.9	48.6	84.3	20.3	81.0	62.9	93.2	87.0	*75.3*
Equity (%) conditional on Equity>0	49.7	34.8	42.5	40.4	43.1	33.9	41.2	41.6	35.4	*43.6*
Money left on the table	102’131	37’355	35’104	53’563	-36’354	51’821	78’107	81’553	99’504	*60’584*
Money left on the table over income	4.2	2.0	0.7	3.9	-0.8	1.3	2.2	2.4	3.6	*2.6*
Borrower’s Age at default	45	43	41	43	44	40	45	45	50	*43*
TTM (years)	15.9	20.9	22.1	27.6	23.9	17.2	14.5	16.9	17.4	*19.6*
Panel B: Loan and borrower characteristics at origination										
LTV (%)	69.9	78.9	79.8	82.5	96.1	90.3	84.1	76.5	101.2	*77.6*
LTI	4.8	7.5	4.1	7.0	4.5	3.1	3.6	3.1	4.3	*5.3*
Loan size (EUR)	121’062	149’245	221’397	96’238	197’278	145’796	127’662	119’439	121’647	*148’131*
Collateral value (EUR)	181’177	193’003	284’785	120’371	211’591	168’732	166’637	160’242	125’457	*196’929*
Loan maturity (years)	24	29	29	35	30	23	21	25	26	*27.0*
Interest only (%)	0.01	0.07	7.60	0.18	43.49	0.32	0.76	25.41	29.63	*4.4*
FRM (%)	22.68	2.65	7.58	0.21	88.48	96.12	82.29	0.20	94.44	*21.6*
Lien *	1.0	1.0	1.0	1.0	1.0	1.0	1.3	1.0	1.0	*1.0*
Household’s income (EUR)	28’418	27’129	59’670	19’495	45’435	55’035	42’611	40’567	30’885	*35’423*
Employed (%)	65.5	66.6	71.8	67.0	77.9	65.9	55.3	56.7	94.4	*67.2*
Selfemployed (%)	21.2	15.6	16.6	15.3	8.4	20.0	22.8	11.7	3.7	*17.7*
Civil/government servant (%)	1.1	3.4	0.0	3.2	0.0	5.3	11.3	0.8	0.0	*2.0*
Pensioner (%)	1.8	2.0	0.1	1.5	1.9	1.0	2.5	0.6	0.0	*1.5*
Number of borrowers *	1.8	1.7	1.6	1.6	1.4	1.6	1.2	1.4	1.4	*1.7*

2.3.3 The distribution of Equity at Default

Table 2.2 provides summary statistics for the sample of defaulted mortgages. Based on our definition above, we classify about 120 thousand loans in default, or about 2% of the overall number of loans. A first striking statistics that emerges from the table is that the average equity at the time of default is a positive and large +27.4%. Even more surprising is the evidence that as many as 75% of defaults, three-fourth of the sample, were above-water. Hence, while much of the extant literature has focused on default in the left tail of the equity distribution, it appears that for Europe above-water defaults are the norm.

The aggregate figures masquerade significant cross-country variability. The average equity at default is highest at +49% for Italy, is right at the average for Belgium, but turns largely negative at -29% for the Netherlands. These differences are partly on account of leverage: the average LTV at origination (Panel B of the table) is nearly 96% in the Netherlands, which makes it more likely that equity becomes negative compared to Italian borrowers with an average LTV of 70%. However, even for the Netherlands we still observe as many as 20% of the defaulted loans having positive equity, with an average equity conditional on being positive of +43.1%, which is very close to the overall +43.6% figure.

How large are these equity percentage figures in monetary terms? In the sixth row of Panel A, we report the money left of the table (in Eur), computed as the difference

Figure 2.4: Equity Distribution at Default

(a) Aggregate distribution

(b) Equity at Default and LTV at origination

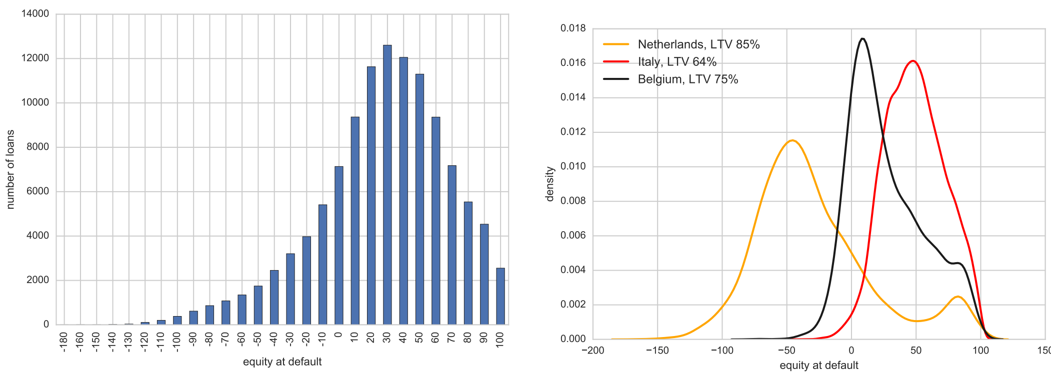

This figures show the distribution of equity at the point when a loan is flagged as defaulted according to the definition in the paper. In panel (a), we plot the histogram of equity at default for the entire sample, where equity is rounded to the nearerst ten. In panel (b), we plot a kernel density estimate for Belgium, Italy and the Netherlands individually. Italy has the lowest average LTV ratio at origination in our sample, while the Netherlands has the highest. Like Belgium, the rest of the countries have an average LTV ratio at origination around 75%. The distributions of the omitted countries lie between those of the Netherlands and Italy.

between the property value and the outstanding balance at the time of default (i.e. the numerator of equation (2.2)). Overall, the value of household's equity stake is about Eur 60'584 higher than the outstanding debt upon default, which corresponds to about 2.6 times the average household's income. If we exclude the Netherlands, for which this ratio is a negative -0.8, borrowers in Ireland leave the smallest fraction of income on the table (0.7) while Italian mortgagers give up at default an equity stake which is on average about four times their annual income.

The remaining part of the table reports loan and borrower characteristics. Compared to those in Table 2.1 for the overall sample, we note that defaulted loans overall feature at origination higher LTV, LTI, maturity, and fraction of ARM, and tend to be issued to borrowers with lower income. However, these conclusions do not hold across all countries.

We report the distribution of equity at default in Figure 2.4. Specifically, Panel A displays the histogram when pooling loans across all countries, whereas Panel B plots a kernel density estimate for three countries: the Netherlands, Belgium, and Italy. The overall distribution is well-behaved with a long left tail, it peaks at about 30% with a significant mass in the 20–50% range, and then reverts towards the 100% maximum. This pattern reassures us that the average values documented above are not driven by outliers but reflect a pervasive phenomenon. When looking into the cross-section, we again note significant differences, with the two extremes being the Netherlands (for which the region of positive equity makes up one third of the distribution) and Italy (which nearly all loans fall into positive territory).

2.4 Empirical Analysis

This section presents our main empirical results. Section 2.4.1 provides some prima facie evidence of the relation between equity and income. We analyze this relation more formally using panel regressions in Section 2.4.2. Section 2.4.3 shows that our findings continue to hold when using alternative property valuations for computing equity.

2.4.1 Bin Scatterplots

Figure 2.5 shows bin scatterplots depicting the relation between equity at default in percentage of the estimated current value of the property and the log of borrower's income at loan origination in the raw data. Since the scatterplots would be very little informative if the relationship was entirely and mechanically driven by differences in the LTV ratio at origination, we only consider loans with a rounded original LTV ratio of 80%, which is by far the most common value. In panel b (panel c) we consider loans that defaulted at negative (positive) values of equity. In every subsample, higher levels of income are associated with lower values of equity at default. This relationship is however most pronounced in the negative equity region.

2.4.2 Main Results

To assess that the positive relationship between equity at default and income is robust to the inclusion of observable loan and borrower characteristics we run the following cross-sectional regression on the subsample of defaulted loans:

$$\begin{aligned} E_{i,t} &= \beta_0 LnIncome_{i,t_0} + \beta_1 LTV_{i,t_0} + \beta_2 LTI_{i,t_0} + \beta_3 Maturity_{i,t_0} + \beta_4 UnempRate_{c,t-12} \\ &\quad + \alpha_l \times \alpha_{t_0} + \epsilon_{i,t} \end{aligned} \tag{2.4}$$

where $E_{i,t}$ is the equity at default of loan i that defaulted in month t (in percentage terms), $LnIncome_{i,o}$ is the household income at origination in logs, LTV_{i,t_0} and LTI_{i,t_0} are respectively the loan-to-value and loan-to-income measured at origination, $Maturity_{i,t_0}$ is the remaining time to maturity of the loan at the time of default, and $Unemp_{c,t-12}$ is the unemployment rate in the region (nuts 2) in the prior year (source is Eurostat).

Finally, our baseline specification incorporates the interaction of county and origination semester fixed effects ($\alpha_l \times \alpha_{t_0}$) to control for local dynamics in house prices and income.

Table 2.3 presents the resulting regression estimates. We uncover a strongly significant and negative relation between equity and income. The full-sample coefficient on log income is -7.14, with a t-statistic of -9.97. This estimate implies that a 10% decrease in household income at origination is accompanied with a 0.71 (percent)

Figure 2.5: Bin Scatterplots

(a) LTV Class = 80

(b) LTV Class = 80, E<0

(c) LTV Class = 80, E>0

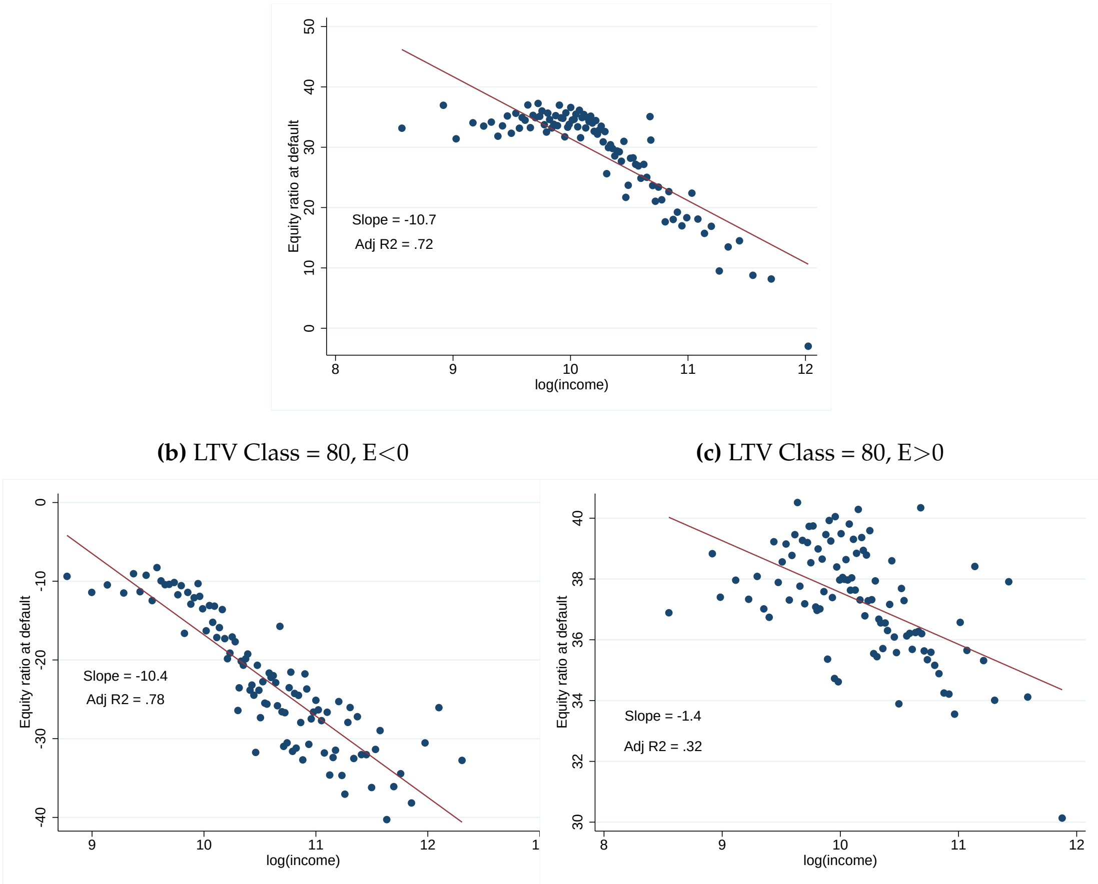

This figure shows bin scatterplots depicting the relation between the equity to value ratio at default and borrower's logincome for loans with a rounded LTV at origination of 80%.

increase in the amount that is left on the table at the time of default. This effect is robust to the inclusion of the other control variables, and in particular LTV and LTI which enter with the expected negative sign – that is, higher leverage and lower levels of affordability are associated with a smaller proportion of equity at the time of default.

To gauge nonlinearities in the relation between equity at default and income, specification (2) adds (log) income squared to the set of regressors. The estimated coefficient is positive and significant, meaning that a decrease in income leads to a more-than-proportional increase in equity. In specification (6), we follow the alternative of including separate dummy variables for borrowers in the top (Q1) to third (Q3) quartile of the income distribution (across all borrowers, defaulted or not) in a given country and month. We observe a similar pattern, namely that everything else constant borrowers at the vertex of the income distribution leave about 6% less equity on the table compared to their low-income peers. To better appreciate the economic magnitude of this effect, consider that borrowers in the bottom quartile of the income distribution who defaulted on their loan have an average income of about

17’000 EUR and an average house valuation of about 126’000 EUR across the panel. Our estimates imply that these borrowers leave about 12’800 EUR additional equity, i.e. about three-fourth of their annual income, compared to high-income households with similar loan characteristics.

In specification (3), we verify that our results remain robust to the inclusion of originator and interest only fixed effects. Specifications (4) and (5) carry the analysis separately for the set of observations with positive and negative equity at default, respectively. In both cases do we observe a negative and significant relation, which is even stronger for above-water borrowers. The result that the threshold level for default is higher for low-income borrowers under negative equity is consistent with earlier literature, but obtains in a setting where the mortgage is with recourse.

In columns (7) and (8), we explore to what extent is the effect of income driven by the fixed versus floating terms of the loan. We thus separately estimate the regression separately for the group of countries with adjustable and fixed mortgages, as defined above. We observe that the loading on income at default is negative and highly statistically significant in both groups. In other words, the decision to leave money on the table at the time of default appears not to be triggered by the interest-rate sensitivity of the loan contract.

Finally, in the last column of the table we use as dependent variable the probability of positive equity defined in equation (2.3). The number of available observations declines due to data requirements to construct the volatility of home prices, as explained in section 2.3.3. Notwithstanding the reduced sample, we note that our findings are robust to using this alternative measure. A 10% decrease in income at origination is accompanied with an increase by about 0.4 percent in the probability of positive equity at default. This effect is statistically significant at the 1% level. LTV and LTI at origination also have the expected negative sign, and are strongly significant.

2.4.3 Robustness Analysis

Since our measure of household equity naturally hinges on the property valuation, we test whether our results are sensitive to the way the latter is computed. While the use of regional house price indices is standard in the literature, see inter alia Bhutta, Dokko, and Shan (2017), our average HPI indices could potentially bias our conclusions if they are not representative of the dynamics of housing prices across the whole income distribution. In particular, if the true real estate prices for cheaper houses plunged more than for expensive houses, our equity at default would be overestimated for low-income borrowers meaning that the coefficient on income in Table 2.3 is biased downward. We address this concern by repeating our analysis on property valuations reported in ED by the RMBS originators, and on foreclosure prices. We find that our main findings continue to hold when computing equity using both these alternative prices.

Table 2.3: Analysis of Equity at Default

This table presents the results for the panel regression of equity at default onto log income and loan characteristics as defined in Table 2.1, and the unemployment rate *Unemp*. Income Q1 to Q3 are dummy variables for borrowers in the top (Q1) to third (Q3) quartile of the income distribution (across all loans, defaulted or not) in a given country and month. t-statistics based on standard errors clustered at the county (nuts 3) level appear in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Dependent Variable	E_t	E_t	E_t	$E_t > 0$	$E_t < 0$	E_t	E_t	E_t	DD_t	E^*_t
Countries	All	All	All	All	All	All	ARM	FRM	All	All
Ln Income	-7.14***	-25.16***	-6.00***	-5.52***	-4.76***		-8.33***	-6.63***	-3.94***	-7.71***
	(-9.97)	(-2.99)	(-8.74)	(-18.17)	(-3.20)		(-6.65)	(-9.27)	(-5.11)	(-7.55)
Ln Income Squared		0.87**								
		(2.05)								
Income Q1 (High)						-8.63***				
						(-10.82)				
Income Q2						-5.95***				
						(-11.01)				
Income Q3						-3.81***				
						(-10.14)				
LTV (%)	-0.94***	-0.94***	-0.93***	-0.60***	-1.01***	-0.95***	-1.08***	-1.06***	-0.62***	-0.80***
	(-28.26)	(-27.88)	(-23.19)	(-43.69)	(-17.13)	(-27.99)	(-22.36)	(-32.06)	(-11.24)	(-30.65)
LTI	-1.24***	-1.32***	-0.93***	-0.93***	-0.59***	-1.01***	-1.17***	-2.22***	-0.90***	-1.83***
	(-8.57)	(-9.86)	(-7.51)	(-9.95)	(-3.05)	(-8.55)	(-6.05)	(-8.30)	(-5.11)	(-6.46)
Maturity	-1.17***	-1.16***	-1.21***	-1.15***	-0.45***	-1.18***	-1.17***	-0.31***	-0.55***	-1.05***
	(-18.20)	(-18.22)	(-17.37)	(-16.02)	(-5.90)	(-17.92)	(-10.70)	(-4.41)	(-7.17)	(-14.91)
Lagged Unemp (%)	-0.60***	-0.60***	-0.58***	-0.21***	-1.30***	-0.60***	-0.89***	-1.47***	-0.58***	-0.18**
	(-4.79)	(-4.84)	(-4.28)	(-6.04)	(-3.05)	(-4.78)	(-4.87)	(-4.34)	(-8.65)	(-2.24)
Country x Orig Semest FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Originator FE	No	No	Yes	No	No	No	No	No	No	No
Interest only FE	No	No	Yes	No	No	No	No	No	No	No
Observations	113,882	113,882	113,824	89,107	24,242	113,882	57,825	13,275	106,511	20,562
R-squared	0.69	0.69	0.72	0.52	0.58	0.69	0.61	0.70	0.53	0.63

Bank-reported property valuations

We make use of updated property valuations that are occasionally submitted in ED by the originators for a subset of the loans. Summary statistics for the sample of available observations are collected in Table 2.4. Overall, we have about 15 million quarter-property observations with a valid price, that is about one-third of the total. For the median loan we observe three valuations, one at loan origination and two in later quarters. The practice of reporting updated property values is, however, quite heterogeneous across countries. In the UK revisions are submitted almost every quarter, while in Spain the property value is rarely reassessed. Unlike valuations at loan origination, updates are usually carried through indexing or AVM. The only exception is Spain, where property inspection is the most common valuation method also in case of updates.

Focusing on this sample of data, we contrast the household's equity that obtains by using alteratively our HPI indices ($\widehat{E}$) or the originator's updated property valuation (E^*). Under the assumption that the latter is more accurately reflecting local conditions and property characteristics, their difference should quantify the extent of measurement error in our estimates. Panel B of Table 2.4 reports the median equity estimation error, defined as $\Delta = \widehat{E} - E^*$. Overall, the error is very small – about 0.3%

pooling all countries, compared with a median equity of about 30% – and is even smaller if we consider properties that have been valuated through a full inspection, whose updates we expect to be most informative.

The fact that the median error is on average nearly zero does not rule out the possibility that it correlates with income distribution. We tackle this concern by regressing the estimation error on log income at origination, controlling for county-origination semester-valuation semester fixed effects. Panel C of the table reports the corresponding slope coefficient and its t-statistic. Two aspects stand out. First, the coefficient is positive, meaning that our measure somewhat overestimates equity for high-income classes. This effect runs in the opposite direction we are concerned about, meaning our results might be slightly conservative. Second, the magnitude of the coefficient is on average 0.17, which is small compared to the estimates of Table 2.3.

As an additional robustness check, we run our main regression model in equation (2.4) on the subsample of defaulted loans for which there is a valuation update in the year in which they became delinquent, using now E^* as our dependent variable. Results are reported in column (11) of Table 2.3. While the number of observations drops considerably, and some of the countries are no longer well represented, we note that the coefficient on log income is even more negative than in our baseline model (moving from -7.25 to -7.71) and remains strongly significant.

Foreclosure prices

A second concern we tackle is that the market prices we use to compute equity may not adequately represent, but rather overestimate the price at which borrowers in financial distress are realistically able to liquidate the collateral to pay back the oustanding principal. In particular, given the illiquid nature of real estate investment, one might argue that the equity at default we document is largely on account of the discount when the borrower is forced to sell it precipitously.

To dispel this concern, we would ideally need to compute equity with respect to the price a distressed household would be able to realize within few months. Unfortunately, this is not feasible with our data. A less ideal, but still useful exercise, is to look at the price at which banks have sold the foreclosed properties. We expect the discount to be larger than what the owner would obtain, bacause (1) foreclosed houses may have been physically damaged during the foreclosure process, (2) financial institutions do not have strong incentives to sell them at the best price, and (3) the properties are usually sold through an auction with a limited number of participants. Therefore, we see the foreclosure price as the lower bound at which a borrower might have been able to sell the house.

In our data, if a foreclosed property has been liquidated, banks are required to report (a lower bound for) the price achieved on the sale. Such price appears for only around 3% of the defaulted loans, which suggests that banks have been and are still

Table 2.4: Analysis using originator's property valuations

This table presents the analysis when using originator's property valuations reported in ED. Panel A presents summary statistics for the sample of available quarter-loan observations. Panel B presents summary statistics, in total or broken down by valuation method, for the equity estimation error defined as the difference between equity computed using HPI indices ($\widehat{E}$) and using the originator's updated property valuation (E^*), that is $\Delta = \widehat{E} - E^*$. Finally, Panel C reports the slope coefficient and its t-statistic in the regression of the equity estimation error Δ on log income, controlling for (county)×(origination semester)×(valuation semester) fixed effects.

	BE	FR	IE	IT	NL	PT	ES	UK	*Total*
Panel A: Summary Statistics									
Obs.	397'487	5'608'658	2'171'834	2'055'360	2'316'714	256'026	50'886	2'566'535	*15'423'500*
Number of valuations per loan									
Mean	2.2	4.5	7.5	4.8	4.1	3.4	1.0	10.8	*4.3*
Median	2.0	4.0	2.0	5.0	2.0	4.0	1.0	14.0	*3*
Valuation Method (%, excluding first valuation)									
Property inspection	3.9	0.0	0.3	7.7	1.5	11.2	75.1	0.1	*2.1*
AVM	0.0	0.0	0.0	19.7	0.0	0.0	0.0	31.4	*7.9*
Indexed	96.1	99.3	99.6	70.9	98.4	80.4	24.9	68.5	*89.7*
Panel B: Equity Estimation Error: Summary statistics									
Median (%)									
All	-4.85	0.78	11.11	0.06	-21.76	10.32	0.00	1.26	*0.27*
Property inspection	1.56		1.36	0.00	-12.42	14.91	0.00	-4.00	*0.00*
AVM				0.00				0.52	*0.36*
Indexed	-5.15	0.81	11.13	0.05	-21.88	11.16	-2.33	1.59	*0.29*
Panel C: Equity Estimation Error: Regression on log income									
Slope	-0.16**	0.07***	1.39***	-0.15	-0.10	-1.49***	-0.16	0.38***	*0.17***
(t-stat)	(-2.08)	(3.50)	(3.84)	(-1.15)	(-0.42)	(-6.07)	(-1.53)	(9.68)	*(2.09)*

reluctant in selling foreclosed properties, most likely to avoid overflowing the market and depressing house prices further. The median reported sale price corresponds to 55% of our estimated market value.[11] As expected, given the large discount, equity at default computed using foreclosure prices drops considerably to a median value of -37%. However, for as many as 33% of the foreclosures the sale price is still larger than the outstanding balance, and the median equity value for abovewater borrowers is as high as 45%.

Given the small number of observations, any analysis using on liquidated properties quickly runs out of statistical power. Nevertheless, Table 2.5 shows that for abovewater borrowers at default, the relationship between equity and income remains negative and significant even when absorbing the variability through several fixed-effect specifications. Only in the case of underwater borrowers we no longer observe a significant relation.

[11] This number is in line with data from Moody's Investors Service 2017 "Prices drop on repossessed Spanish homes, as liquidation gathers pace" that the average sale price of repossessed Spanish homes sold in 2015 and 2016 also amounted to roughly 45% of the properties indexed valuation, reflecting an illiquid repossessed property market.

Table 2.5: Equity using Foreclosure Prices

This table presents the results for the panel regression of equity at default onto log income and controls, when equity at default is constructed using foreclosure prices. t-statistics based on standard errors clustered at the county (nuts 3) level appear in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Dependent Variable	(1) $\tilde{E}_t$	(2) $\tilde{E}_t$	(3) $\tilde{E}_t$	(4) $\tilde{E}_t > 0$	(5) $\tilde{E}_t < 0$
Ln Income	-20.58***	-20.67**	-21.43**	-6.65**	7.65
	(-2.18)	(-2.36)	(-2.53)	(-2.29)	(1.16)
Controls	Yes	Yes	Yes	Yes	Yes
Origination Semester FE	No	Yes	Yes	Yes	Yes
Default Semester FE	No	No	Yes	No	No
Observations	5,401	5,401	5,400	1,395	4,004
R-squared	0.07	0.09	0.11	0.28	0.08

2.5 Default sensitivity to equity shocks

Section 2.4 presents results from the analysis conditional on the borrower having defaulted on the loan. The evidence that low income borrowers default at higher levels of equity is only suggestive of a higher sensitivity to the value of collateral. In this Section, we take a panel-data approach to ask the following question directly: Conditional on the same current level of equity, are low-income borrowers more likely to default than higher income borrowers when facing an housing shock?

To this end, we collapse the data to yearly frequency and we estimate the linear probability model in equation (2.5) on the 40 million loan-year observations in ED

$$D_{i,t} = \gamma' \Delta HPI_{i,t} \times IncQ_{i,t_0} + \beta E_{i,t-1} + \delta' W + \alpha_t + \alpha_l + \alpha_{t_0} + \epsilon_{i,t}, \tag{2.5}$$

where $D_{i,t}$ is a dummy that equals 1 for a loan i that defaults in year t, $\Delta HPI_{i,t}$ is the growth in house prices in the region of loan i between year $t-1$ and t, $IncQ_{i,t_0}$ is a vector of income quartile dummies and $E_{i,t-1}$ is the equity ratio in the previous year. The vector W collects control variables. In particular, we control for LTI at origination, loan term and lagged NUTS2 unemployment rate. This approach allows us to include a wider array of fixed effects. In our baseline specification we include county, year and loan origination year fixed effects to control for differences in average default rates across these dimensions. Results are reported in Table 2.6. Column (2) combines county times year fixed effects and column (3) only includes country and year fixed effects. We cluster the standard errors by region to account for any potential within-region dependencies over time. In column (4) we also estimate a probit model which produces similar results.

Our main interest is on the coefficients on the interaction terms. The results confirm that defaults are sensitive to changes in house prices and that this sensitivity

Table 2.6: Analysis of probability of default

This table presents the results for the panel regression of default rates. The dependent variable is a dummy that equals one if the loan defaulted, and zero otherwise. The dependent variables are defined as in Table 2.3. *HPI Ret* is the return to the HPI in that county-semester. The regression is estimated with a linear probability model in columns (1)-(3), and as a probit in column (4). t-statistics based on standard errors clustered at the county (nuts 3) level appear in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively. For the probit specification, the R-squared is the pseudo R-squared.

	(1) Linear	(2) Linear	(3) Linear	(4) Probit
$E_{i,t-1}$	-0.414***	-0.420***	-0.389***	-0.445***
	(-6.545)	(-6.389)	(-9.013)	(-6.388)
Income Q1 (High)	-0.051***	-0.056***	-0.067***	0.000
	(-5.013)	(-5.565)	(-6.928)	(0.012)
Income Q2	-0.043***	-0.050***	-0.055***	0.015
	(-4.999)	(-5.631)	(-6.296)	(0.780)
Income Q3	-0.016***	-0.019***	-0.022***	0.034***
	(-3.027)	(-3.308)	(-3.916)	(2.757)
HPI Ret	-0.694***	-0.711***	-0.427***	-1.249***
	(-7.427)	(-7.251)	(-4.139)	(-10.806)
HPI Ret$\times$ Income Q3	0.194***	0.200***	0.254***	0.229***
	(2.971)	(3.052)	(3.681)	(4.055)
HPI Ret$\times$ Income Q2	0.344***	0.344***	0.425***	0.347***
	(3.963)	(3.925)	(4.964)	(4.598)
HPI Ret$\times$ Income Q1	0.415***	0.401***	0.433***	0.262***
	(3.992)	(3.827)	(4.120)	(2.725)
TTM	-0.001	-0.001	-0.001**	-0.011***
	(-1.188)	(-1.445)	(-2.033)	(-5.223)
Unemp	2.460***	1.555***	1.038***	2.060***
	(4.901)	(5.776)	(4.961)	(8.557)
LTI	0.013***	0.013***	0.010***	0.029***
	(7.341)	(7.238)	(6.525)	(7.255)
Constant				-2.865***
				(-47.770)
Observations	49,124,528	49,124,528	49,124,525	49,124,528
R-squared	0.006	0.006	0.010	0.0428
County FE	Yes	-	-	-
OY FE	Yes	Yes	-	-
Year FE	Yes	Yes	-	-
Country FE	-	Yes	-	-
Country x OY x Year	-	-	Yes	-

is inversely related to income. As expected, the estimated effect of lagged equity is negative and significant indicating that the probability of default is decreasing in equity. The lower part of the table shows that the control variables matter and that their effect is of the expected sign.

2.6 Supply side channel

So far, we have considered mortgage default as the outcome of a one-sided decision by the borrower. However, borrowers struggling to meet their mortgage obligations may try to negotiate with the lender possible solutions (reinstatement, repayment plan, forbearance or loan modification) to prevent foreclosure before becoming deeply delinquent. By lowering the financial burden on borrowers, such loan renegotiations are expected to have a significant impact on default probabilities.

The eligibility of borrowers strongly depends on the probability of the loan to become current again, conditional on the modification. To the extent that income at origination might be related to borrowers experiencing temporary rather than persistent income shocks, our results might be capturing the effect of (a lack of) access to measures of financial relief rather than borrower-level considerations on the benefits of default.

To dig into this alternative explanation, we look at whether financial constraints at the lender side affect the level of equity at default for different income levels. We capture the extent of credit constraints on the supply side by sorting originators based on their ratio of non-performing-loans (NPL). We use the detail of ED data to track an originator's loan portfolio performance over time and compute the value of loans in distress. Specifically, in a given quarter, we compute NPL as the ratio between the cumulative sum of the balance of non-performing-loans and the cumulative sum of the balance of the originated loans. If the loans in ED are representative of an originator's overall loan portfolio, we can regard NPL as a proxy for the originator's financial constraints. We then create quartiles based on the overall NPL distribution[12] and create dummy variables Q1-Q4 for a given originator, which equal 1 if the originator's most recent NPL falls in the top (Q1) to bottom (Q4) quartile of NPL and zero otherwise. Notice that, in this way, our measure of lender financial constraints varies both across counties and in the time-series.

If lenders are less inclined to renegotiate loans for lower income borrowers, we expect the relationship between equity at default and income to be amplified when lenders experience financial distress, since originators that experienced higher losses on their assets have less leeway to modify loans. We test this hypothesis by rerunning our analysis of defaults by interacting log income with dummy variables Q1-Q3, so that the non-interacted coefficient captures the effect for mortgagers whose loan is with an originator with a top performing portfolio (i.e. least constrained). Table 2.7 collects the corresponding estimates, under various fixed-effect specification. We see that the interaction terms are almost always monotonically decreasing. That is, low-income borrowers leave more equity at default, but this effect is about twice as large when the originator is in the top quartile of NPL. For example, in the first column, the baseline coefficient of -4.11 becomes -8.73 when the originator has

[12] The ratio of non-performing loans in the lowest quartile is 0.9%, compared to 8.2% is the higher quartile.

Table 2.7: Credit supply channel

This table presents the results for the panel regression of equity at default onto log income and controls, when log income is interacted with dummy variables for originators' NPL quartiles. In a given quarter, we compute NPL as the ratio between the cumulative sum of the balance of non-performing-loans and the cumulative sum of the balance of the originated loans. We then create quartiles based on the overall NPL distribution and create dummy variables Q1-Q4 for a given originator, which equal 1 if the originator's most recent NPL falls in the top (Q1) to bottom (Q4) quartile of NPL and zero otherwise. t-statistics based on standard errors clustered at the county (nuts 3) level appear in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

	(1)	(2)	(3)	(4)	(5)
Dependent Variables	E_t	E_t	E_t	E_t	E_t
Ln Income	-4.11***	-3.21***	-4.09***	-2.49***	-3.01***
	(-5.56)	(-3.68)	(-5.24)	(-2.82)	(-3.40)
Ln Income * NPL Q3	-3.62**	-4.51***	-3.52***	-4.81***	-4.55***
	(-2.52)	(-2.84)	(-2.72)	(-3.29)	(-2.84)
Ln Income * NPL Q2	-4.28***	-5.54***	-4.25***	-4.52***	-5.45***
	(-2.63)	(-3.05)	(-2.97)	(-2.75)	(-2.98)
Ln Income * NPL Q1 (High)	-4.62**	-5.66***	-4.10***	-5.30***	-5.55***
	(-2.48)	(-3.09)	(-2.97)	(-3.28)	(-2.98)
NPL Dummies	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes
Country FE	Yes	No	No	No	No
County x Origination Semester FE	No	Yes	Yes	Yes	Yes
Default Quarter FE	No	No	Yes	No	No
Bank FE	No	No	No	Yes	No
Payment Method FE	No	No	No	No	Yes
Observations	112,234	111,640	111,634	111,593	111,640
R-squared	0.57	0.70	0.70	0.73	0.70

faced severe losses in its loan portfolio. This pattern is consistent with a credit supply story, whereby the most constrained originators curtail credit to lower-income borrowers. At the same time, however, the coefficient on log income in the first row remains large and statistically significant. This finding is robust both to the inclusion of quarter as well as lender fixed effects, showing that the effect exists both across lenders and across time.

We interpret this evidence as suggesting that, everything else constant, lending frictions matter in determining the level of equity at default, and partly explain why low income borrowers default at higher levels of equity. However, the fact that the coefficient on the reference group, for which financial constraints on the originator's side are lowest, remains negative and significant indicates that our results are not merely capturing a credit-supply effect.

2.7 Conclusions

Mostly due to lack of data, we know very little about mortgage performance in Europe. There is a large literature that investigates the drivers of mortgage defaults using data from the US mortgage market, especially since the recent subprime crisis. Though, there are many institutional differences between mortgage markets in the euro area and the United States, in particular the fact that in Europe nearly all mortgages are full recourse.

The financial literature mostly focuses on the decision to default in the context of non-recourse mortgages. The main predictions are that (1) falling house prices increase the probability of default by increasing the value of the embedded option to default (strategic default), (2) negative equity is a necessary but not sufficient condition of default and (3) the threshold level of equity is increasing in borrowing constraints.

Under recoursability, borrowers are liable thorugh personal assets of the remaining balance upon default. First, recoursability is expected to discourage borrowers from defaulting when equity is negative and in general should decrease the sensitivity of default to negative equity, as shown empirically in Ghent and Kudlyak, 2011.

However, in this paper we posit that even without the embedded option of non-recourse mortgage contracts, there still exists a threshold level of equity that would trigger default. If this is the case, everything else constant, we expect that the threat of lenders recourse would be mostly feared by borrowers with higher marginal utility of consumption, higher assets to be seized, or both. We test this prediction on our data by relating equity at default with income and various borrowers characteristics.

We show that a large fraction of defaults in Europe happen when borrowers are above water: the majority of borrowers who default on their loan occurs when the value of their collateral would be in principle enough to repay the debt, with estimated equity at default of 25% on average. Consistent with the threat of recoursability being greater for borrowers with a high marginal utility of current consumption, we find that equity at default is significantly negatively related with the household's income at origination. This result persists even after controlling for the loan-to-value and loan-to-income ratios, and when scaling the equity by price volatility. We also find evidence that a tightening in credit conditions reinforces the relation between equity at default and income and negative shocks to house prices have a larger impact on the default probability of low-income borrowers, everything else controlled for.

References

Acharya, Viral V et al. (2017). "The Anatomy of the Transmission of Macroprudential Policies: Evidence from Ireland". In: *16th International Conference on Credit Risk Evaluation, Interest Rates, Growth, and Regulation, September*, pp. 28–29.

Adelino, Manuel, Antoinette Schoar, and Felipe Severino (2016). “Loan originations and defaults in the mortgage crisis: The role of the middle class”. In: *The Review of Financial Studies* 29.7, pp. 1635–1670.

Bajari, Patrick, Chenghuan Sean Chu, and Minjung Park (2008). *An empirical model of subprime mortgage default from 2000 to 2007*. Tech. rep. National Bureau of Economic Research.

Bhutta, Neil, Jane Dokko, and Hui Shan (2017). “Consumer Ruthlessness and Mortgage Default during the 2007 to 2009 Housing Bust”. In: *The Journal of Finance* 72.6, pp. 2433–2466.

Bhutta, Neil, Hui Shan, and Jane Dokko (2010). “The depth of negative equity and mortgage default decisions”. In:

Campbell, John Y (2012). “Mortgage market design”. In: *Review of finance* 17.1, pp. 1–33.

Campbell, John Y and Joao F Cocco (2015). “A model of mortgage default”. In: *The Journal of Finance* 70.4, pp. 1495–1554.

Elul, Ronel et al. (2010). “What “triggers” mortgage default?” In: *American Economic Review* 100.2, pp. 490–94.

Ertan, Aytekin, Maria Loumioti, and Regina Wittenberg-Moerman (2017). “Enhancing Loan Quality Through Transparency: Evidence from the European Central Bank Loan Level Reporting Initiative”. In: *Journal of Accounting Research* 55.4, pp. 877–918.

Feldstein, Martin (2008). “How to Help People Whose Home Values Are Underwater”. In: *Wall Street Journal*. URL: https://www.wsj.com/articles/SB122697004441035727.

Flodén, Martin et al. (2017). “Household debt and monetary policy: Revealing the cash-flow channel”. In:

Foote, Christopher L, Kristopher Gerardi, and Paul S Willen (2008). “Negative equity and foreclosure: Theory and evidence”. In: *Journal of urban economics* 64.2, pp. 234–245.

Geanakoplos, John (2014). “Leverage, default, and forgiveness: Lessons from the American and European crises”. In: *Journal of Macroeconomics* 39, pp. 313–333.

Gete, Pedro and Franco Zecchetto (2018). “Mortgage Design and Slow Recoveries: The Role of Recourse and Default”. In: *Available at SSRN 3082409*.

Ghent, Andra C and Marianna Kudlyak (2011). “Recourse and residential mortgage default: evidence from US states”. In: *The Review of Financial Studies* 24.9, pp. 3139–3186.

Kau, James B., Donald C. Keenan, and Taewon Kim (1994). “Default Probabilities for Mortgages”. In: *Journal of Urban Economics* 35.3, pp. 278 –296. ISSN: 0094-1190. DOI: https://doi.org/10.1006/juec.1994.1017.

Mian, Atif and Amir Sufi (2009). “The consequences of mortgage credit expansion: Evidence from the US mortgage default crisis”. In: *The Quarterly Journal of Economics* 124.4, pp. 1449–1496.

Van Bekkum, Sjoerd, Marc Gabarro, and Rustom M Irani (2017). "Does a Larger Menu Increase Appetite? Collateral Eligibility and Credit Supply". In: *The Review of Financial Studies*.

3 Quantitative Easing and Equity Prices: Evidence from the ETF Program of the Bank of Japan

3.1 Introduction

With policy interest rates constrained at the zero lower bound, many central banks around the world have resorted to unconventional monetary policy tools. Within the range of unconventional measures, large-scale asset purchase (LSAP) programs have attracted particular attention because of their large size and thus their impact on central banks' balance sheets. The massive expansion of both the assets and liabilities of central banks exposes them to considerable risks and raises questions about the consequences of a potential exit from QE.

There is considerable evidence that central banks' asset purchases can have an economically significant impact on yields in the targeted markets, which has likely motivated central banks to continue these purchases over the past decade (D'Amico and King, 2013; Eser and Schwaab, 2016; Gagnon et al., 2010; Hamilton and Wu, 2012; Krishnamurthy and Vissing-Jorgensen, 2011; Krishnamurthy and Vissing-Jorgensen, 2013; Neely, 2010; Swanson, 2011). However, despite the widespread use of LSAP programs, the debate is still ongoing with regard to the mechanisms linking asset purchases to asset prices and the persistence of the impact. Unlike policy rate targeting, asset purchases are explicit decisions on quantities and are designed to have a noticeable impact on market prices. Even though the idea of *easing through quantity* relies on the view that large purchases by the central bank reduce assets' risk premia, there is still no clear theoretical foundation for how and under which conditions this is expected to work. In general, the relationship between the outstanding quantity of an asset and its price is not yet well understood.

Since 2013, the Bank of Japan (BoJ) has been engaging in what they have named *Quantitative and Qualitative Easing* (QQE) program as an attempt to fight against deflation. As part of its broader QQE agenda, the BoJ has been vigorously increasing its domestic equity holdings through purchases of index-linked ETFs. By the end of 2016, the BoJ owned approximately ¥14 trillion worth of TOPIX and Nikkei ETFs, which corresponds to more than 2.5% of the total market capitalization. This unprecedented equity operation has the declared objective of lowering the risk premia of asset prices and reducing the cost of equity capital of Japanese companies (BoJ, 2013).

The BoJ is the first central bank to engage in purchases of domestic equities as part of its QE agenda. This intervention represents a unique laboratory to shed light on the long-standing debate on the elasticity of long-term demand curves for stocks, as well as a unique opportunity to test how QE impacts equity prices and its implications for market efficiency. In this paper we study its impact on the cross-section of stock prices. We propose a novel empirical strategy to identify and quantify the price impact of the change in assets' supply through QE. This provides new evidence on the price elasticity of long-run demand curves.

The literature on the effectiveness of QE has proposed several channels through which central banks can affect prices. A natural explanation is provided by the so-called "portfolio-balance" channel, first discussed by Brunner and Meltzer (1973) and Frankel (1985) and Tobin (1969). According to this channel, when the central bank buys a particular asset, it reduces the amount held by the private sector, effectively changing the risk composition of the aggregate portfolio held by investors. For this to be an equilibrium, prices need to adjust to ensure market clearing.

In this paper, we first propose a simple structural asset pricing model that generalizes the idea of the portfolio balance channel to the case of equities.[1] The key implication of portfolio rebalancing that we derive from the model is that the change in systematic risk of each stock is determined by: (i) the entire vector of central bank purchases and (ii) the covariance matrix of stocks' cash flows.

We then bring the model to the data in a standard event-study framework, exploiting two specific events in which the BoJ announced major expansions of its ETF purchases. On October 31, 2014, the BoJ announced a three-fold increase in the purchase of ETFs and on July 29, 2016, it communicated a further doubling of the budget amount.

We document that both announcements produced a highly heterogeneous response of equity prices at the company level. Figure 3.1 plots the cumulative returns of two portfolios following the 2014 announcement, formed by ranking stocks on the price impact predicted by the model. The divergence in returns is statistically and economically significant. Results from cross-sectional regressions show that the variation in event returns in the cross-section is consistent with the change in the marginal contribution of each stock to the risk of the aggregate portfolio held by private investors, as predicted by the portfolio-balance channel.

Looking at longer-horizon returns we find no evidence of reversal over a one-year window after both policy announcements, which supports the main time-series prediction of the model. We estimate the long-term net effect of the portfolio-balance channel at about 22 basis points increase in market value per trillion Yen employed

[1] It is easy to show that the duration channel discussed in the literature is a special case of our model when all securities in the economy are exposed to a single source of risk.

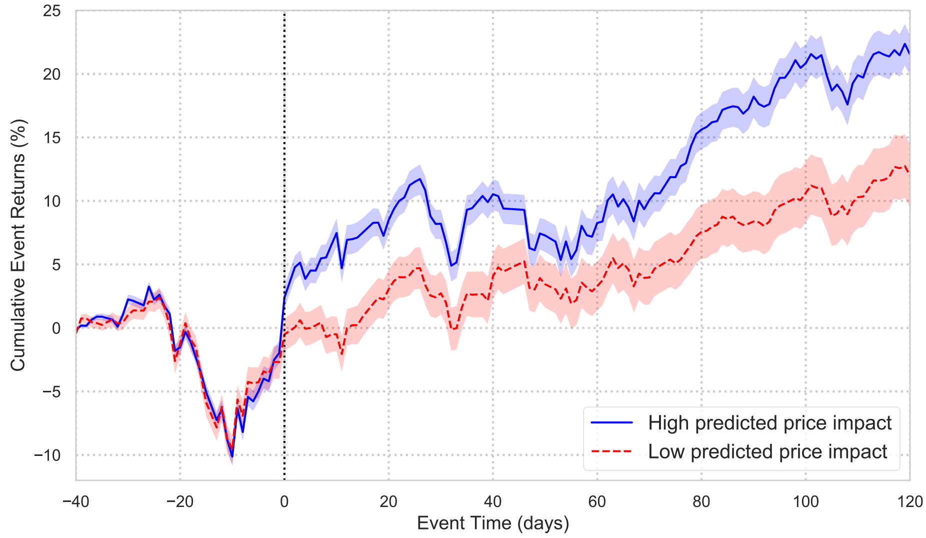

Figure 3.1: Heterogeneity of Event Returns. This figure shows the time series of the cumulative returns around the BoJ announcement of October 2014, of two portfolios of firms ranked by the model predicted returns. The blue line is the average for first quartile of the distribution (firms with the highest predicted price impact), while the red dashed line corresponds to the average for the last quartile (firms with the lowest predicted price impact). Bands represent bootstrapped 95% confidence intervals.

in the program. Given a total equity market capitalization of about ¥500 trillion, this implies an elasticity close to 1, so that each yen invested translates into an increase in total market valuation of roughly one yen. Our estimate is in line with those provded by (Shleifer, 1986a) and (Petajisto, 2011), who find an elasticity of 1 and 0.84, respectivetly, using additions to the S&P 500 index. The analysis based on Dutch auction repurchases of (Bagwell, 1992) also results into a relatively close price elasticity of 1.65 Other authors, instead, find flatter demand curves with price elasticity ranging from the value of 8.24 estimated by (Wurgler and Zhuravskaya, 2002) to that of 10.5 by (Kaul, Mehrotra, and Morck, 2000).

The two expansions of the policy budget provide us with an ideal natural experiment to examine the net effect of a long-lasting change in supply on prices for three reasons. First, the purchase schedule of the central bank is exogenous to firms' fundamentals in the cross-section. Second, unlike asset purchases by the Federal Reserve, the program of the BoJ affects the supply of each security according to an ex ante well-defined purchase schedule. Third, since roughly half of the capital of the central bank is allocated according to the weights of the price-weighted Nikkei 225 index, the purchases produce variation in the cross-section of supply shocks relative to market capitalization that is as good as random. In general, the identification of the impact of LSAPs on asset prices is a challenging task. The intervention of the BoJ provides us with an empirical framework that mitigates endogeneity concerns and improves the identification of the net (short-run and long-run) effect of a change in supply, which crucially relies on the exogeneity of the shock.

The non-fundamental nature of the Nikkei weighting system has already been exploited in Greenwood (2005) and Greenwood (2007) to establish a causal relationship between uninformed demand shocks and prices in the context of a large redefinition of the Nikkei 225 index. A major difference between the LSAP setting and the one of index redefinitions lies in the nature of the supply shocks. As the central bank buys assets, it is effectively transferring a portion of fundamental risk from the private sector to its balance sheet, and holds it for an arguably long period of time. This is at least conceptually different from an index redefinition event, in which securities merely change hands from active investors to index funds. The central bank can be thought of as a buy-and-hold long-term investor whose portfolio holdings are not marked-to-market. Its long-term commitment to the policy induces a long-lasting change to supply, making our setting better suited than index redefinitions to identify long-run price effects due to movements along investors' long-term demand curves.

The model that we propose extends the theoretical framework in Greenwood, 2005 to account for this difference in setting. As in Greenwood, 2005, we consider an economy with multiple assets in finite supply and a CARA-utility representative agent that maximizes her wealth in each period. We introduce quantitative easing in the form of an exogenous shock to the supply of assets, which is first announced and then gradually carried out over a given policy horizon. The agent correctly understands that the QE program will affect the market-clearing portfolio in each future period, which determines the new vector of equilibrium risk premia. Crucially, we extend the model to an infinite horizon to relax the assumption that uncertainty is resolved at a terminal date, which mechanically drives the reversal in Greenwood, 2005. In our model, prices adjust to the change in supply to reflect the new risk composition of the aggregate portfolio held by the representative agent. Unless the central bank is expected to unwind its position, this implies that we should not observe a reversal at any horizon. The fact that we observe a persistent effect in the data is consistent with this prediction of the model.

In the data, not only we find no evidence of a reversal of the initial jump in prices, but non-trivial abnormal returns are still observed one year after the announcements. Even though the policy is carried out gradually, the total size of the intervention is revealed to the market in advance. Market efficiency requires that today's prices reflect expectations about future returns, hence they should also reflect expected future changes in the supply of assets. In the model, the observed post-event drift can arise because of two reasons. First, since purchases are scattered over various dates, the model predicts that prices continue to adjust also after the announcement due to the decrease in the residual duration of the program over time. However, for realistic levels of the risk free rate, we show that this effect is expected to be quantitatively small relative to the initial price jumps. A more pronounced drift arises when we allow the representative agent to believe that the central bank will deviate from the announced purchase target. The model produces a sluggish price reaction similar to the one observed in the data when expectations about the size of the purchase program are assumed to increase over time, consistently with investors underreacting to the announcement as well as with learning about additional purchase programs in the future.[2]

We address the concern that (part of) the observed price impact and its persistence might be explained by repeated price pressure rather than a portfolio-balance channel. Large trades from the BoJ may give rise to order imbalances, thus pushing prices upwards on purchase days. Such mispricings are expected to be shortly lived in efficient markets. However, arbitrageurs may refrain from trading if the central bank is expected to buy again soon, thus failing to bring prices back to their fundamental value. This would imply a persistent price effect of the program arising from the flow of the purchases rather than the change in the supply of assets, a channel quite distinct from portfolio balance. As in D'Amico and King, 2013, we will refer to this effect as the *flow effect* of the program. Even though the difference may seem subtle, disentangling between these two channels has important implications. First, the two channels lead to different conclusions about the elasticity of long-run demand curves for stocks. Second, they imply different consequences of a potential exit from QE. In particular, if QE is mainly effective through repeated price pressure, a slow-down or a suspension of the purchases would cause a sharp drop in prices. On the contrary, in our model of the portfolio-balance channel, it is not the flow into the balance sheet of the central bank that keeps prices up, but its accumulated size. Therefore, suspending the purchases should have a more limited effect on prices. We exploit both the cross-sectional and time series variation in purchase volumes to identify and quantify the flow effect of the policy in the spirit of Eser and Schwaab (2016). We then re-estimate the cross-sectional portfolio-balance channel effect using returns net of the flow-induced component. We find that price pressure effects are positive and persistent. However, this channel might explain at most a minor fraction (between 12% and 23% depending on the specification) of the estimated portfolio balance effect.

Overall, our empirical analysis confirms the concerns raised by the financial press that the intervention of the BoJ might be inducing price distortions due to the deviation of the purchase schedule from market weights. We document a significantly heterogeneous effect of the policy both at company and industry level. A modification of the QQE has the potential to address this problem. Theoretically, the only way to achieve a cross-sectionally homogeneous shift in risk premia is for the BoJ to hold each stock proportionally to the company market capitalization. At the moment, however, still roughly a quarter of the BoJ capital is allocated to the price-weighted Nikkei index.

The rest of the paper is organized as follows. Section 3.2 describes the ETF purchase program of the BoJ. Section 3.3 reviews the relevant literature. Section 3.4 presents

[2] Beliefs are exogenous in our model and evolve deterministically over time. Extending the model to a setting where beliefs are endogenous is definitely interesting, but beyond the scope of the paper.

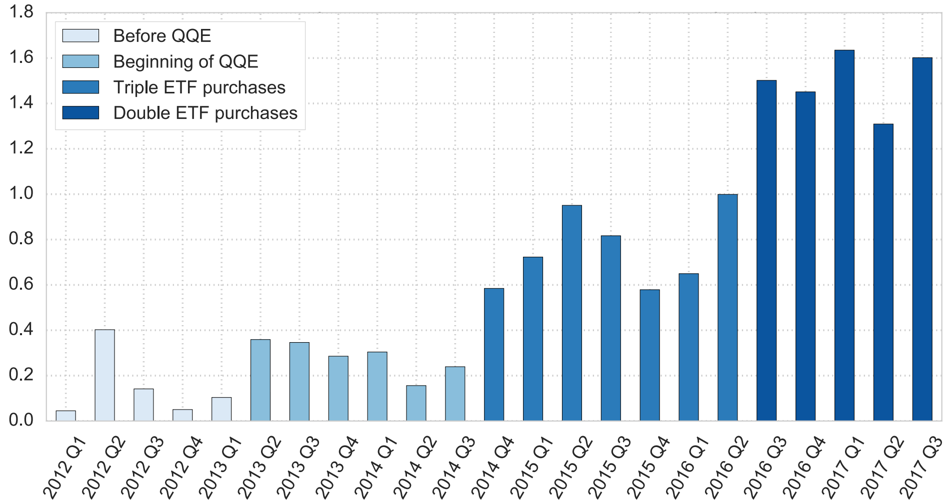

Figure 3.2: Quarterly ETF Purchases of the Bank of Japan in trillion yen. Changes in the bar color indicate changes in the policy target purchase amounts. In the first phase the target was set to ¥1 trillion per year, in the second phase it was tripled to ¥3 trillion and in the third phase it was additionally doubled to ¥6 trillion. Data is from the BoJ website.

the model. Section 3.5 describes the data, the empirical strategy and estimation procedures. Section 3.6 presents our main empirical findings. Section 3.7 considers the flow effect of direct purchases and evaluates its relative importance with respect to the portfolio balance effect. Section 3.8 discusses policy implications and Section 3.9 concludes.

3.2 The ETF Program of the BoJ

As part of the "Quantitative and Qualitative monetary Easing" (QQE) introduced on April 4, 2013, the BoJ embarked on a large-scale asset purchase (LSAP) program committing itself to buy large quantities of broad market equity ETFs, with the declared view of lowering risk premia (BoJ, 2013). The policy budget was initially set at ¥1 trillion per year (roughly US$ 10 billion). On two occasions, the BoJ announced a sharp expansion of the target amount: on October 31, 2014, the Bank communicated that the annual mark was tripled to ¥3 trillion, and was again doubled on July 29, 2016, to ¥6 trillion. The policy changes are clearly visible along the time series of monthly ETF purchases by the bank, as shown in Figure 3.2. The time series of aggregate ETF purchases is publicly available at daily frequency on the BoJ website starting from December 2010.

Its holdings accumulated rapidly, and by the end of 2016, the BoJ owned more than ¥14 trillion worth of ETFs. This corresponds to 2.5% of the total capitalization of the First Section of the Tokyo Stock Exchange (TSE), and around 3% of the Japanese GDP. The share of BoJ holdings to aggregate Assets Under Management (AUM) of targeted ETFs has grown from almost zero to more than 70% since the beginning of the program; this is even more remarkable if we consider that the ETF industry in Japan almost tripled in value between 2013 and 2016. In terms of size, the ETF program is comparable to the annual aggregate net flows into or out of the Japanese equity fund industry and therefore economically relevant.[3]

The purchase program targets two types of ETFs: those tracking the Tokyo Stock Price Index (TOPIX) and those replicating the return of the Nikkei 225 Stock Average.[4] At inception of the program, the money allocated to each ETF was set to be proportional to its assets under management (AUM). The ratio of the aggregate AUM of ETFs tracking the TOPIX Index and those of ETFs tracking the Nikkei 225 Index is roughly 1 to 1.2. This approximately translates into half of the capital flowing into Nikkei ETFs and half into TOPIX ETFs. In turn, this then maps into a demand shock at the stock level that depends on each company's weight in the corresponding index.

The TOPIX is a value-weighted index tracking the roughly 2000 companies listed on the First Section of the TSE, while the Nikkei 225 is a *price-weighted* index of 225 TOPIX companies representative of the Japanese stock market. The constituents of the Nikkei index are typically large blue-chip companies that account for roughly two-thirds of the market capitalization of the TSE First Section on aggregate. The Nikkei 225 is the most widely traded equity benchmark in Japan.

The weighting system of the two indices implies that the BoJ allocates only half of its budget to companies proportionally to their market value. The remaining half of the budget flows instead to the Nikkei constituents proportionally to their price, not accounting for the number of shares outstanding, thus producing mis-allocation relative to market capitalization. Under market efficiency, the market value of a company should reflect all available fundamental information. The dispersion of the ratio between price weights and value weights is therefore expected to be unrelated to firms fundamentals. The relative under-weighting in the BoJ portfolio is clearly more severe for companies not included in the Nikkei index. However, there is a high degree of heterogeneity in the allocation of capital across Nikkei companies as well. This is clear from Figure A.1 in Appendix A, where we plot the distribution of the log of the ratio between the weight in the Nikkei and the weight in the TOPIX for Nikkei companies to measure the cross-sectional dispersion of the resulting allocation at the stock level.

Given the unusual weights of the BoJ purchase schedule, a sudden expansion of the policy budget produces a natural experiment where stocks are hit by an uninformed demand shock that is highly heterogeneous in the cross-section and orthogonal to firms fundamentals after controlling for market capitalization. In this paper, we

[3]Over the past 10 years, the average net flows into equity funds in Japan was roughly ¥3 trillion in absolute value per year. Data are from the Thomson Reuters Lipper Global Fund Flows database.

[4]On November 19, 2014, the BoJ started buying also ETFs tracking the JPX-Nikkei 400 Index. This approximately corresponds to 43% of the purchases flowing to ETFs tracking the TOPIX, 53% to ETFs tracking the Nikkei 225 and the remaining 4% to ETFs tracking the JPX Nikkei 400. For simplicity, in the empirical analysis we round the share of both TOPIX and Nikkei ETFs to 50%, neglecting the JPX Nikkei 400. This simplification does not affect the results of our analysis.

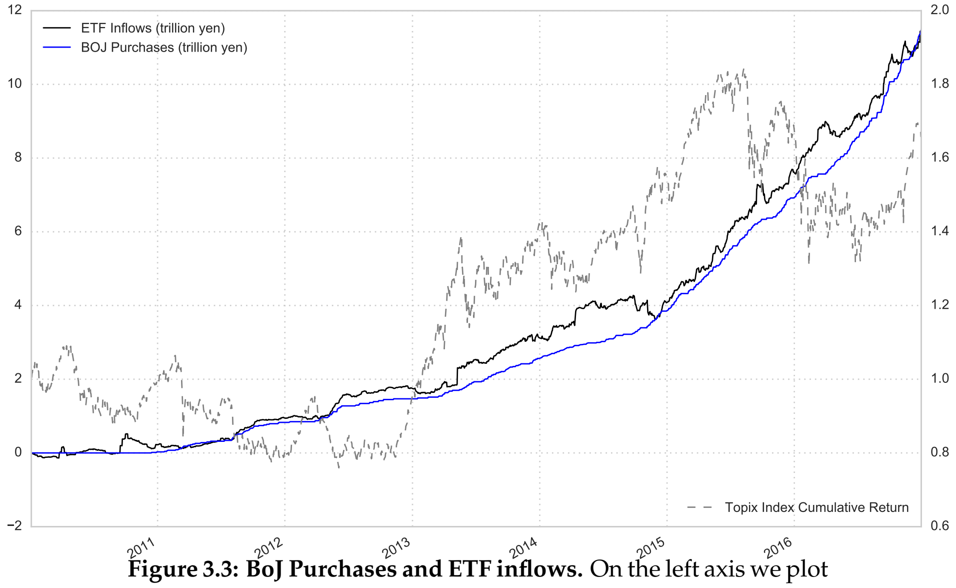

Figure 3.3: BoJ Purchases and ETF inflows. On the left axis we plot the daily cumulative purchases of ETFs by the BoJ (blue line) and the estimated daily cumulative inflows into ETFs tracking either the Nikkei or the TOPIX index (black line). Both are in trillion yen. On the right axis the figure shows the cumulative return from 2010 of the TOPIX index (gray dashed line).

exploit the exogenous variation in the cross-section of supply shocks to identify the causal impact of the purchase program on equity prices. We rely on a simple asset pricing model to argue that the deviation from a value-weighted allocation allows us to isolate the portfolio-balance channel of the policy impact. Section 3.5.2 discusses the identification strategy in detail.

Overall, the portfolio of the BoJ ends up deviating significantly from the allocation that market capitalization would dictate. To illustrate the extent of this distortion, take three companies with fairly similar market capitalization and therefore similar TOPIX weights (between 0.45% and 1% in 2014): Canon, Fast Retailing and Nintendo. Canon and Fast Retailing are both among the Nikkei constituents, though with very different weights, namely around 1.2% versus 9.5%, respectively. Nintendo, on the contrary, is not included in the Nikkei index. It follows that the BoJ allocates to Fast Retailing 4 times more capital than to Canon, and 19 times more than to Nintendo. The effects of the departure from a value-weighted allocation are reflected in the indirect ownership that the BoJ accumulated over time. According to estimates by the Financial Times, through its purchases the central bank has indirectly become the largest shareholder in a quarter of TOPIX stocks. In Table 3.1 we report the ten stocks with the highest estimated indirect ownership share by the BoJ.

We argue that purchases of ETFs by the BoJ translate into supply shocks at the individual stock level. This is a consequence of the creation-redemption mechanism in the ETF primary market and the physical replication of the underlying basket. When demand exceeds supply in the ETF secondary market, new shares of ETF are issued to keep the ETF price close to its NAV. In the case of physical ETFs, creation requires

Company Name	BOJ Share (%)	BOJ Flow (bn JPY)	Market Cap (bn JPY)	Nikkei weight (%)
Mitsumi Electric Co Ltd	10.3	5.8	56.1	0.17
Advantest	8.9	27.5	309.1	0.63
Fast Retailing	8.7	336.4	3854.7	9.17
Taiyo Yuden	7.8	10.0	129.0	0.32
Toho Zinc	7.7	3.3	43.2	0.09
Tdk Corporation	7.4	71.0	959.0	1.41
Konami Holding	7.2	37.6	524.5	0.65
Trend Micro	7.0	36.2	514.9	0.90
Comsys Holding	6.6	18.3	275.7	0.39
Nissan Chem In	6.2	30.6	489.7	0.53
Average	6.1	3.8	255.6	0.44
Median	5.9	0.2	43.8	0.20

Table 3.1: BoJ indirect shareholdings. Summary statistics on indirect ownership by the BoJ for the ten companies with the highest BoJ share. BoJ Flow are the cumulative compounded BoJ purchases at company level since the beginning of QQE and Market Cap is the company's market capitalization. BoJ Share is the ratio of BoJ Flow and Market Cap. Average and median values are calculated over the universe of TOPIX firms. The values in the first three columns are as of August 31, 2016. The last column reports the average company weight in the Nikkei 225 index over the study period. Notice that the ten companies with the highest BoJ share have all positive weights in the Nikkei 225 index.

the physical purchase of the basket of securities that composes the tracked index, for a value equal to the creation unit. Securities are then held by the ETF sponsor on behalf of the owner of the ETF shares, who now bears the associated risk. ETF creation thus reduces the quantity of assets available for trading in the underlying market. This mechanism is visualized in Figure A.2 in Appendix A. Given the direct correspondence, in the rest of the paper we will consider ETF purchases by the central bank equivalent to an intervention in the underlying equity market.

We can infer whether central bank purchases triggered creation of new ETF shares from data on ETFs AUM. We first obtain the list of the ETFs listed on the Tokyo Stock Exchange (TSE) that track either the Nikkei or the TOPIX index from the website of the Japan Exchange Group (JPX). We then get daily data on AUM for each ETF from Bloomberg. We estimate inflows simply as the difference between the actual increase in AUM and the increase in AUM due to the return on the index that the ETF is tracking. Figure 3.3 plots the time-series of ETF inflows versus the amount purchased by the BoJ. It is apparent that the flows into these ETFs are almost completely due to the asset purchase program. In turn, this implies that the purchases by the BoJ have consistently triggered creation of new ETF shares.

It must be noted that the bias towards Nikkei companies did not go unnoticed among practitioners and the BoJ was frequently accused by the financial press of distorting the market. In response to the criticism, on September 21, 2016, the BoJ amended the terms and conditions of the program and announced it will change the maximum amount of each ETF to be purchased. Since October 2016, the BoJ allocates ¥2.7 trillion a year (US$ 26.4 billion) to TOPIX ETFs, while the remaining ¥3 trillion are spread out between the TOPIX, the Nikkei 225 and the JPX-Nikkei Index 400. For the Nikkei-ETFs this means a drop from 55% to about 25% of the annual purchases by the BoJ, which brings the allocation of the flows closer to what market capitalization would justify. Yet, the accumulated balance sheet of the BoJ remains tilted away from a value-weighted allocation.

3.3 Related Literature

"Extraordinary times call for extraordinary measures", stated the Chairman of the Federal Reserve Ben Bernanke in 2009 (Bernanke, 2009). Since then, a number of central banks around the world have adopted unconventional monetary policy tools and most of them have been trying to support asset prices through LSAPs in order to boost economic activity in the face of severe dislocations in financial markets. With actual data on the implementation of LSAPs becoming available, a large body of academic research has investigated their impact on financial markets and the real economy.

Most of the work on the impact of QE on market prices relies on evidence from purchases of government bonds by the Fed, the ECB or the BOE, and usually shows a significant impact on yields (Buraschi and Whelan, 2015; D'Amico and King, 2013; Eser and Schwaab, 2016; Gagnon et al., 2010; Hamilton and Wu, 2012; Joyce et al., 2012; Krishnamurthy and Vissing-Jorgensen, 2011; Krishnamurthy and Vissing-Jorgensen, 2013; Neely, 2010; Swanson, 2011). There is however little empirical evidence on the large-scale purchases of the BoJ. Perhaps closest to our paper is Ueda (2013), who looks at the time series of LSAP announcements by the BoJ and finds a positive correlation with the TOPIX index and the yen-dollar exchange rate. The BoJ is the first central bank to purchase domestic equities as part of its QQE agenda, and, to the best of our knowledge, this paper is the first to study this program in depth and to analyze its impact on the cross-section of stock prices.

Although there is general agreement that LSAPs do indeed affect prices, there is less consensus regarding the channels through which these policies work. A standard explanation in the literature is the so-called portfolio balance channel (Brunner and Meltzer, 1973; Frankel, 1985; Tobin, 1969). According to this channel, when the central bank buys a particular asset, it reduces the amount held by private investors, effectively forcing them into a different portfolio. For this to be an equilibrium, prices need to adjust to ensure market clearing. In particular, through this channel asset purchases are expected to push up the price of the target asset and of its substitutes, implying that demand curves are downward sloping. Some papers find that the observed price impact is consistent or partially consistent with portfolio balance explanations (e.g. D'Amico and King (2013), Gagnon et al. (2010), and Joyce et al. (2011))[5]. However, the portfolio balance channel of monetary policy is subject of debate, in part because standard asset pricing models do not generally allow exogenous changes in the supply of a security to affect its price. For instance, Miles and Schanz (2014) argue that LSAPs by central banks since 2008 had significant effects because markets were dysfunctional and that in normal times portfolio-balance effects would be weak.

The question whether demand curves for stocks slope down has a long tradition in the asset pricing literature. The empirical evidence so far mostly comes from event studies around index redefinitions and fire sales by institutional investors (Coval and Stafford, 2007; Greenwood, 2005; Harris and Gurel, 1986; Hau, Massa, and Peress, 2009; Mitchell, Pulvino, and Stafford, 2004; Petajisto, 2009; Scholes, 1972; Shleifer, 1986b; Schnitzler, 2016). The general finding is that large non-fundamental trades have a significant but temporary price impact, even though there is considerably heterogeneous evidence on the speed and the extent of reversal. The standard interpretation is that limits to arbitrage can justify temporary deviations from fundamental value: under market efficiency, uninformed shocks cannot have a long-lasting impact on prices.[6]

Quantitative easing provides an ideal laboratory in which to test asset pricing theories such as the long-held belief of flat demand curves for stocks. However, from the success of QE in pushing up prices alone, one cannot conclude much about the elasticity of demand curves. A large number of papers show that LSAPs by central banks have effects beyond those due to portfolio balance, and provide evidence of alternative transmission channels that are consistent with flat demand curves. For the case of purchases of long-term bonds, Krishnamurthy and Vissing-Jorgensen (2011) provide compelling empirical evidence that the so-called signalling channel explains a significant fraction of the drop in bond yields observed after the Federal Reserve's QE announcements. The idea behind this channel is discussed in Eggertsson and Woodford, 2004, who claim that financial markets may interpret LSAPs as signals about the central bank's intention to keep interest rates low, thus influencing long-term yields through investors' expectations about the future path of interest rates. Other papers attribute the beneficial effect of the Fed's MBS purchases on risk premia during the financial crises to a capital constraints channel motivated by the distress in the financial intermediary sector (Curdia and Woodford, 2011; He and Krishnamurthy, 2013).

In general, the identification of the impact of market interventions through a specific channel is a challenging task. Our paper contributes to this literature proposing a new identification strategy for the transmission channel of monetary policy and

[5]Vayanos and Vila (2009) try to reconcile the predictions of the portfolio balance channel with the observed lack of spillovers across maturities, building on market segmentation and preferred-habitat theories as proposed by Culbertson, 1957 and Modigliani and Sutch, 1966.

[6]The traditional view in finance is that, in a frictionless world, a simple expansion of the balance sheet of the central bank should have no effect. This neutrality result is formalized in Eggertsson and Woodford, 2004 and crucially relies on the assumption of a rational infinitely lived agent with no credit restrictions, who sees no difference between its own assets and those held by the central bank.

providing new insights on the elasticity of demand curves. Moreover, the results of the empirical literature suggest that the specific workings of LSAPs depend on the asset purchased and the economic conditions under which these purchases take place. We complement the existing evidence by documenting the effects of the ETF program by the BoJ, a unique case in which a central bank is targeting the equity market.

3.4 The Model

In this section we develop a theoretical framework to describe the portfolio balance channel as the transmission mechanism from LSAP to asset prices. The idea is that asset purchases shift part of the fundamental risk from the market to the balance sheet of the central bank. Because the premium demanded for a given security is proportional to its marginal risk contribution to the aggregate portfolio held by the representative agent, the price effect of the monetary intervention is proportional to the implied change in this quantity. Therefore, the net effect on asset prices through this channel is not simply proportional to the purchased amounts, but it crucially depends on the correlation structure of firms fundamentals.

Our model features the central bank only in reduced form, in the sense that the policy rule is exogenous. We also assume that asset purchases are deterministic. This assumption holds also when we allow investors to believe that the central bank will deviate from the announced purchase target. With no policy uncertainty, asset purchases do not represent a source of risk that has to be priced in equilibrium. Moreover, we assume firms fundamentals to be neutral with respect to monetary policy, excluding the possibility that asset purchases affect market prices through the change in future investment opportunities. We make these choices to keep the model simple and to focus on the direct effect of supply on prices. These assumptions also allow us to restrict our attention to the covariance-stationary equilibrium of the model, which immediately follows once we assume covariance-stationary dividends. The limitation is that the model abstracts from potential additional channels related to uncertainty about future supply and endogenous responses of firms.

3.4.1 Model Setup

Consider an economy with n risky assets in fixed supply $Q = (Q^1, \dots, Q^n)$, paying dividends in every time period. The dividend $D_{i,t}$ paid at time t is

$$D_{i,t} = D_{i,0} + \sum_{s=1}^{t} \varepsilon_{i,s}, \qquad \forall i \in 1, \dots, n \tag{3.1}$$

where each $\varepsilon_{i,t}$ is revealed at time t. The fundamental innovations $\varepsilon_{i,t}$ are modelled as zero-mean jointly normal random variables, iid over time.

The representative agent optimally chooses her time-t demand N_t to maximize her next period expected utility, subject to a standard budget constraint

$$\max_N \; \mathrm{E}_t\left(-\exp(-\gamma W_{t+1})\right) \tag{3.2}$$

$$\text{s.t.} \qquad W_{t+1} = W_t(1+r) + N_t'(p_{t+1} + D_{t+1} - p_t(1+r)) \tag{3.3}$$

where W_t is the total wealth, N_t' denotes the transpose of the vector N_t and γ the aggregate risk-aversion. At date $t = 1$ the central bank announces share purchases described by the vector $q = (q^1, \ldots, q^n)$, distributed over M periods after the announcement. We refer to M as the policy horizon. Let q_t denote the vector of cumulative purchases by the central bank up to date t. One can think of q_t as the active side of the balance sheet of the central bank at any time t.

We assume, first, that $q_t = tq$ for $t = 1, \ldots, M$ and, second, that $q_t = Mq$ for $t > M$. The first assumption implies that in our model the central bank's balance sheet evolves deterministically and grows linearly over time. Assuming non-stochastic asset purchases allows us to abstract from policy uncertainty as a priced risk factor and to focus on how QE affects prices through the change in supply [7]. The second assumption implies that the central bank never unwinds its position nor engages in further purchases beyond horizon M. This assumption might be restrictive once we go to the data since the BoJ never announced such a stringent commitment. Still, given that the BoJ position have not been unwound (and neither announced to be so) over the window of our empirical analysis, we believe it to be a reasonable benchmark.

The realized demand shocks negatively affect the net supply of assets in each period. Setting $Q_0 = Q$ yields

$$Q_t = Q - q_t \tag{3.4}$$

Asset purchases by the central bank affect the quantity at which the equity market clears given the equilibrium condition $N_t = Q_t$. Notice that equation (3.4) also implies that the quantity of assets available to the market can only change through purchases of the central bank. This excludes the possibility for companies to respond endogenously to changes in prices by issuing new stocks or buying back those outstanding.

The central bank buys the vector q of securities in exchange for cash. We assume that the representative agent invests the proceeds in the risk-free asset and, since risk-free returns are uncorrelated with those of Japanese equities, omitting the risk-free asset from the model does not change the predicted policy impact on stock prices. This assumption may be interpreted as a form of market segmentation, in that the representative agent cannot re-invest the proceeds in assets outside the Japanese equity universe. We discuss the implications of this assumption in Section 3.6.4. The

[7] The assumption that the central bank spreads its purchases equally over the policy horizon is instead innocuous. Relaxing this assumption does not improve the economic intuition and only adds technical complexity to the model.

fact that we do not model other asset classes that equity investors might hold in their portfolios does not affect the model predictions even in case of non-zero correlation with equities. It is easy to show that including securities that are not targeted by the asset purchase program has no effect on the predicted price impact on stock prices. Stock purchases will spillover to correlated asset classes, but in this paper we are not interested in these effects.

Appendix B shows that the pricing equation of the covariance-stationary equilibrium, in matrix notation, is given by

$$p_t = \frac{1}{r}\left(D_t - \gamma V \Omega_t\right) \tag{3.5}$$

where $V \equiv \text{Var}_t(p_{t+1} + D_{t+1})$ is the stationary covariance matrix of asset returns and Ω_t is the vector of time-t expected future asset supply, properly discounted by time, defined as

$$\Omega_t \equiv \frac{r}{1+r} \sum_{i=0}^{\infty} \frac{\mathrm{E}_t[Q_{t+i}]}{(1+r)^i} \tag{3.6}$$

Notice that the covariance matrix V is not time-varying because the supplies of each asset are fixed and the schedule of purchases by the central bank is deterministic, so there is no uncertainty on future shocks to the asset supply. Equation (3.6) shows that at any point in time prices reflect the path of future asset supply. Given the time-discounting, today's prices are less sensitive to quantities further into the future.

Staring at the vector of risk premia $\gamma V \Omega_t$ in the pricing equation (3.5) one can see that, for each stock, priced risk is an increasing function of the stock's covariance with the market portfolio and the risk aversion parameter γ. The vector $V\Omega_t$ admits an interpretation very similar to the CAPM beta and should be thought of as a measure of systematic risk.[8] This is easier to see in the absence of monetary policy shocks, in which case $V\Omega_t$ reduces to VQ.

By plugging equation (3.4) into equation (3.5) it is easy to see how the portfolio balance mechanism works in the model. Asset purchases change the amount of each security in the market clearing portfolio. This affects systematic risk and in turn prices. Notice that this change in systematic risk is fully consistent with our assumption of a constant covariance matrix V, since what determines systematic risk is the product $V\Omega_t$, and central bank purchases affect only the latter term in the model.

Let's now turn to how the representative agent builds expectations about future asset supply, before and after the purchase program is announced. These expectations enter the pricing equation (3.5) and so determine the impact of the policy. At time $t > 0$, the representative agent's expectation about the quantity in period

[8] While the model is written in price changes, market betas are usually defined in terms of returns. In Appendix C we derive an expression of systematic risk that determines expected returns in the model. While the notation becomes messier, the intuition carries through.

$h = t+1, \ldots, M$ is given by

$$E_t[Q_h] = E_t[Q_0 - q_h] = Q - E_t[q_h] \tag{3.7}$$

We assume that the central bank intervention is fully unexpected at $t = 0$ before the announcement, i.e. $E_t[q_h] = 0$ for every $t \leq 0$ and $h \geq t$. At each period $t \geq 1$ after the announcement date, we allow the investor to believe that the central bank will deviate from its purchase target. We restrict to a family of investors beliefs parametrized by a time-varying scalar λ_t. Formally, let $\lambda_t \geq 0$ be real numbers such that

$$\mathrm{E}_t\left[Q_h\right] = Q - \lambda_t q_h, \qquad t \geq 1 \tag{3.8}$$

The parameter λ_t is assumed to change over time in a deterministic fashion. The path of λ_t determines how the representative agent updates her beliefs regarding the size of the purchase program. We first solve the model for a general mapping $t \mapsto \lambda_t$ and then we present, in the next section, results for the special case $\lambda_t \equiv 1$, in which the pricing equation takes a simpler form that better conveys the intuition for the portfolio balance channel. Imposing $\lambda_t \equiv 1$ is equivalent to assuming that the representative agent expects the central bank to commit to the announced target exactly. This implies that she also expects the central bank to never unwind its positions and to never engage in additional purchase programs in the future. Changing the parameter λ_t allows us to study how deviations from this benchmark case impact the effect of the policy.

As shown in Appendix B, the model's pricing equation predicts price changes given by

$$p_t - p_{t-1} = \frac{1}{r}\left(\varepsilon_t + \gamma\xi(t)Vq\right) \tag{3.9}$$

where the function $\xi(t)$ is defined piece-wise as follows

$$\xi(t) = \begin{cases} 0 & \text{if} \quad t \leq 0 \ \text{ or } \ t > M \\ \lambda_1(M - \varphi(1)) & \text{if} \quad t = 1 \\ \Delta\lambda_t M - (\lambda_t\varphi(t) - \lambda_{t-1}\varphi(t-1)) & \text{if} \quad 1 < t \leq M \end{cases} \tag{3.10}$$

and $\varphi(t) < M$, defined in Appendix B, is a deterministic function of time representing the residual duration of the purchase program.

In the first part of equation (3.10), $\xi(t) = 0$ implies that both before the announcement ($t \leq 0$) and after the purchase program has been completely carried out ($t > M$), price changes only reflect shocks to dividends and are therefore unpredictable. The functional form of $\xi(t)$ in the second and third pieces of the domain determine event ($t = 1$) and post-event price changes ($1 < t \leq M$), respectively. Even when future supply changes are fully predictable and the average path of future prices can be perfectly anticipated ($\lambda_t \equiv 1$), the shock to supply is impounded into stock prices immediately after the policy announcement only up to the term $\varphi(1)$. Prices will then continue to adjust in the following days. The reason why prices do not fully adjust on the event day is that until purchases are actually realized at future dates, the representative agent bears dividend risk and requires a compensation for it. Consistent with this intuition, $\varphi(t)$ is decreasing in t and increasing in M. So, even though $\xi(t)$ drives predictable post-event price changes, these are fully consistent with market efficiency and do not represent an arbitrage opportunity.

The relative magnitude of the initial price reaction and the subsequent adjustments depend on the level and the dynamics of λ_t. More specifically, the price jump at $t = 1$ is increasing in the initial expectation of future supply λ_1 since prices are effectively responding to a purchase program of size $\lambda_1 Mq$. Post-event price changes are then linked to the time series evolution of λ_t. An increasing λ_t over time means that the agent is revising upward her expectations about the size of the program. One can think of different reasons for why this might happen. For example, the agent may not immediately believe that the central bank will commit to the full size of the program and thus update her expectations only once she observes the purchases actually being carried out. Or, she may start believing over time that the central bank will engage in additional purchases beyond the announced policy horizon M. Similarly, a decreasing λ_t means that the agent revises downward the expected size of the program, either because she starts to believe that the central bank will not complete the announced program or that it will unwind the portfolio soon after. In the Internet Appendix we show simulations of the price dynamics implied by different functional forms for λ_t.

Even though we mainly think of λ_t as controlling the agent's beliefs on the central bank actions conditional on time-t information, this reduced-form suits a number of non-mutually exclusive interpretations. For instance, as in Barberis and Thaler, 2003, the slow reaction may be due to the bounded rationality of agents who fail to correctly process the consequences of the BoJ announced program.

Benchmark Case

In this section we focus on the special case where $\lambda_t \equiv 1$, which implies that expected and realized purchases are the same at any point in time

$$E_t[q_h] = q_h = h\,q \tag{3.11}$$

It follows directly from equation (3.9) that the price adjustment at $t = 1$ is given by

$$p_1 - p_0 = \frac{1}{r}\left(\varepsilon_1 + \gamma V(Mq - \varphi(1)q)\right) \tag{3.12}$$

Ignoring fundamental innovations, equation (3.12) predicts a positive price jump of magnitude $\gamma V(Mq - \varphi(1)q)$. This swing in prices is due to the fact that the policy is unexpected at $t = 0$, but it is impounded into prices as soon as it is revealed.

In the following periods ($t \geq 1$), price changes are instead given by

$$p_{t+1} - p_t = \frac{1}{r}\left(\varepsilon_{t+1} - \gamma V(\varphi(t+1) - \varphi(t))q\right), \qquad t = 1, \ldots M \tag{3.13}$$

Equation (3.13) shows that price changes in the post-announcement period include a non-stochastic component $\gamma V(\varphi(t+1) - \varphi(t))q$ which accounts for the time delay between the announcement of the supply shocks and their realizations. Since the cross-sectional distribution of these predictable price adjustments is always parallel to and of the same sign as the initial price impact, they add up to create a propagation (drift) of the initial cross-sectional effect.

3.4.2 Testable Predictions

In this section we derive testable predictions from the model. To make these predictions more suitable to be tested in the data, we state them in terms of returns. In order to go from the expressions in price changes derived in Section 3.4.1 to predictions about returns, we first need to introduce some new notation. We define u as the vector of yen amount purchased by the BoJ of each security, so that

$$u_i \equiv p_{i,t}\, q_i, \qquad \forall i \in 1\ldots, n \tag{3.14}$$

where q_i is the number of shares purchased of stock i and $p_{i,t}$ the stock price at time t. Then, we define Σ to be the stationary covariance matrix of stock returns, i.e.

$$\Sigma_{i,j} \equiv \mathrm{Cov}(R_i, R_j), \qquad \forall i, j \in 1\ldots, n \tag{3.15}$$

where R_i is the daily percentage return of stock i.

Dividing equation (3.5) by p_{t-1} leads to the following two propositions about event returns and post-event returns. Proofs are in Appendix B.

Proposition 1 (Event returns). *The vector of returns $R_1 = (p_1 - p_0)/p_0$ on the announcement day is positively related to the vector $\pi \equiv \Sigma u$ in the cross-section.*

Proposition 2 (Post-event returns). *Assume $\Delta\lambda_{t+1} = \lambda_{t+1} - \lambda_t \geq 0$. Then the vector of post-event returns R_{t+1} is positively related to $\pi = \Sigma u$ in the cross-section for every $t = 1, \ldots, M$. Moreover, the vector of expected cumulative returns is given by*

$$\sum_{s=1}^{t} \mathrm{E}[R_s] = \theta_t \pi \tag{3.16}$$

where $\theta_t = \frac{\gamma}{r}\sum_{s=1}^{t} \xi(s)$ is a positive and increasing function of t, which follows from the definition of $\xi(t)$ in equation (3.10).

Proposition 1 states that through the portfolio balance channel, the policy announcement leads to abnormal event returns proportional to the change in systematic risk captured by the vector $\pi = \Sigma u$. Notice that if the central bank were to buy stocks proportionally to their market weight, abnormal event returns would be proportional to the product of Σ and the vector of market capitalizations, i.e. the vector of each stock's covariance with the market portfolio. Proposition 1 therefore implies that an exogenous shock to supply parallel to the market portfolio would cause price adjustments proportional to market betas. At the same time, it also implies that shocks to supply that are orthogonal to market capitalization produce abnormal returns orthogonal to market betas. This prediction is key to identify the effect of the policy shock from the cross-section of realized event returns in the empirical part of the paper.

As summarized in Proposition 2, the model predicts post-event returns in the same direction of event returns, i.e. proportional to π, until the purchase target is met at $t = M$. This generates a post-event drift, whose magnitude depends both on the value of the risk free rate and the beliefs dynamics parametrized by λ_t. In the Internet Appendix, we show analytically and from model simulations that for realistic value of the risk free rate and λ_t constant, this drift is small. The model produces a more pronounced drift under the assumption that the representative agent revises her expectations on the size of the program over time ($\Delta\lambda_{t+1} > 0$).

From Proposition 2 it follows that a permanent change in the supply of assets generates a permanent change of risk premia, and hence of prices. Unless the central bank unwinds its positions, prices will not revert to the pre-event level.[9] By stating that changes in supply can have long-lasting impacts on prices, the proposition implies downward sloping demand curves for stocks through the portfolio balance mechanism.

3.5 Data and Empirical Methodology

3.5.1 Data Sources

From Compustat Global we collect stock-level data on daily returns, volumes and shares outstanding for the roughly 2000 stocks of the TOPIX universe for the period 1990-2016. Daily returns and volume data for the TOPIX index as well as the monthly time-series of TOPIX and Nikkei 225 index weights for every stock in our sample are obtained from Thomson Reuters Datastream. The USD/JPY exchange rate is from Japan Macro Advisors Inc. The time-series of ETF purchases by the BoJ is publicly available at daily frequency on its website.

[9]Notice that a reversal would be observed as soon as investors update the expected path of future supply to include a sale of the portfolio of the central bank ($\Delta\lambda_t < 0$). Also, we would observe a reversal if the central bank was to surprise the market by ceasing the purchases before reaching the expected target. Since we do not provide any empirical evidence of an exit from LSAP, we do not formalize this scenario into a proposition.

3.5.2 Identification Strategy

To test the model predictions from Proposition 1 and Proposition 2 we estimate the following cross-sectional regression at different horizons H around the two policy announcements made by the BoJ

$$R^H_{i,e} = \alpha_e + \beta^H_e \pi_{i,e} + \delta'_e W_{i,e} + \eta_{i,e} \tag{3.17}$$

where R^H_i is the cumulative return of stock i computed over H days from the event day and W is a matrix of stock-level covariates. The estimation of the vector π is described in the next section. All variables are event specific and therefore indexed by the subscript $e \in (2014, 2016)$. Regression coefficients are also indexed by the event because we estimate the model separately for the two announcements.

The coefficient of interest β^H measures the portfolio balance effect of the policy and is identified from the cross-sectional heterogeneity of the model-implied change in systematic risk π. Notice that β^H has a similar interpretation as the coefficient on the interaction term in a diff-in-diff estimation where π measures the intensity of the treatment.

Following Proposition 1, if stock returns respond to the exogenous shock to supply through the mechanism described in the model, we expect $\hat{\beta}^H$ to be positive and significant at short horizons. Proposition 2 implies a positive and significant coefficient at any horizon H. We therefore look at $\hat{\beta}^H$ estimated from a regression of cumulative returns over longer horizons (one month, three months, six months and one year) on π. Estimating a positive $\hat{\beta}^H$ at short horizons followed by a lower $\hat{\beta}^H$ at longer horizons would indicate that the initial event return is, at least partially, reversed after some time. Such evidence would be inconsistent with the portfolio balance channel and would rather suggest a temporary price pressure story, where arbitrageurs with limited capital need some time to absorb the demand shock coming from the central bank. Proposition 2 also implies that $\hat{\beta}^H$ should be found to be weakly increasing in H. An increasing $\hat{\beta}^H$ indicates that the divergence in the cross-section of returns in the direction of the vector π not only does not vanish, but it becomes larger with time.

The identifying assumption behind this strategy is that there is no transmission mechanism of monetary policy other than portfolio balance that would affect prices proportionally to π. If the central bank was buying according to market weights, π would be parallel to market betas and this claim would be hard to make, as we know from the literature that asset purchases can affect stock prices through multiple channels, possibly in proportion to their exposure to market risk. However, since the BoJ is tilting its purchases away from market capitalization, Proposition 1 predicts that the shock should leave a characteristic footprint in the cross-section of returns and abnormal returns, which makes the assumption more likely to be true.

3.5.3 Variable Construction and Summary Statistics

This section presents the data and defines the empirical proxies for the vector u of expected purchases and of the covariance matrix Σ of asset returns, defined in Section 3.4.2.

To be conservative, in Section 3.6 we test the model predictions separately for the two policy announcements of the BoJ. All variables are therefore calculated or estimated twice, in order to have two sets of variables, one for each event. For all estimated variables in our analysis we use an estimation window of one year, ending two trading weeks before each BoJ announcement.

Expected Purchases

In the guidelines to the LSAP program, the BoJ states that it would spread its purchases among index-tracking ETFs proportionally to the aggregate AUM of each ETF. In practice, this roughly corresponds to a 50-50 allocation of capital between TOPIX and Nikkei ETFs.

When we go to the data, we assume that this allocation rule not only holds on aggregate over the policy horizon, but also each time the central bank makes a purchase. Under this assumption, the vector u of purchases by the BoJ (in yen) can then be expressed as

$$u_i = Tw_{i,T} + Nw_{i,N} \tag{3.18}$$

where T and N indicate the amount of BoJ capital allocated to TOPIX and Nikkei ETFs, respectively, and $w_{i,T}$, $w_{i,N}$ are the weight of stock i in the TOPIX and Nikkei indices.

Since $T \cong N$ in the current purchase program, for the empirical analysis we compute the vector u simply as $w_{i,T} + w_{i,N}$. The vector u proxied in this way still encodes the cross-sectional variation in purchases at the heart of our identification strategy. Given that index weights are time varying, for each event we take $w_{i,T}$ and $w_{i,N}$ to be the index weights as of the end of the month preceding the announcement.

Figure A.4 in Appendix A shows in the top row of each panel, the cross-sectional distribution of stock weights in the TOPIX and the Nikkei 225 index in the month before the event. The top-right panel plots the cross-sectional distribution of the resulting stock-level weights in the BoJ purchase vector $(w_{i,T} + w_{i,N})$. The percentile plots in logarithmic scale clearly show that variation in weights across stocks is substantial.

Covariance Matrix

We estimate the variance-covariance matrix Σ of stock returns using daily returns data from Compustat Global. Because the cross-sectional dimension of our data is larger than the sample size, the sample covariance matrix of returns is a poor estimator of Σ. We therefore use the shrinkage method proposed by Ledoit and Wolf, 2004 to obtain a well-conditioned and more accurate estimator, which also ensures that the resulting matrix is always positive definite.

In the model described in Section 3.4 returns are driven only by fundamentals innovations and changes in supply. However, when we go to the data, this assumption may not hold. We are especially concerned about the impact on the returns moments of other monetary policy announcements during the estimation window. To address this concern, we look at stock returns net of market returns and we estimate Σ as the cross-sectional covariance of the fitted residuals $\hat{e}_{i,t}$ from a simple market model specified as

$$R_{i,t} = \alpha_i + \beta_i^{mkt} R_{mkt,t} + e_{i,t} \tag{3.19}$$

where $R_{i,t}$ are daily returns of stock i and $R_{mkt,t}$ is the return on the TOPIX Index used as proxy for the market portfolio. As reported in Table A.2 of the Appendix, our results are robust to estimating Σ on raw returns rather than abnormal returns.

Control Variables

We estimate stocks' sensitivities to changes in the exchange rate by running the following regression separately for each stock i

$$R_{i,t} = \alpha_i + \beta_i^{mkt} R_{mkt,t} + \beta_i^F F_t + e_{i,t} \tag{3.20}$$

Here F_t is the daily percentage change in the exchange rate from US Dollar to Japanese Yen. Estimation results for market and Forex betas are reported in Table A.1 and Figure A.4 in Appendix A. In the bottom rows we plot the cross-sectional distributions of pre-event market betas and of Forex betas together with companies' market capitalization. Table A.1 in Appendix A presents a break down of the summary statistics by Nikkei and non-Nikkei companies.

3.6 Empirical Results

In this section we test the empirical predictions of the model described in Section 3.4. Our results show that the ETF program of the BoJ had a significant impact on stock prices and that both the cross-sectional and time-series patterns of the price effect are consistent with a portfolio balance channel. We first perform event studies around the two BoJ announcements and show that the observed price impact is positively related at the stock-level with our ex-ante measure of systematic risk change. We then look at the effect over different horizons and conclude that the impact of the policy is persistent and increasing in time. We propose a simple back-of-the-envelope calculation to quantify the net aggregate portfolio balance effect of the implemented program. We estimate a 22 basis points increase in aggregate market valuation per trillion Yen invested, which corresponds to a unitary price elasticity.

3.6.1 Event Study

Proposition 1 states that we should observe a positive relationship between each security abnormal event return and the change in its marginal contribution to the risk of the aggregate portfolio.

As a preliminary test of this relationship we rank stocks in the TOPIX universe by the predicted abnormal event return $\pi_i = (\Sigma u)_i$ into four equally-weighted portfolios. Figure 3.4 presents cumulative returns of the low and high π portfolios. Plots on the left show the event returns around the first policy change in 2014 (when the target purchase amount of ETFs was tripled), while those on the right present the effect of the second change in 2016 (when the target was doubled further). We consider raw returns and abnormal returns based on two versions of the market model with different proxies for the market portfolio, the TOPIX index and an equally-weighted index, respectively. The reported bands represent bootstrapped 95% confidence intervals.

Each plot shows a sizeable and highly significant spread between the returns of high and low π firms opening after the two announcements. While for the 2014 event the reaction seems to be slightly anticipated, in 2016 the effect is delayed by a couple of days. Overall, the pattern of abnormal returns is similar for the two events, with the performance of the high π portfolio being significantly higher than that of the low π portfolio. There is no sign of reversal over 30 days after the announcement, and rather the gap between the two groups appears to increase over time. This preliminary evidence is consistent with both predictions of the model.

3.6.2 Cross-Sectional Regressions

One might be concerned that, by sorting on π, we are implicitly ranking stocks based on firms' characteristics such as size, export share or market beta, which might explain the heterogeneous response to the announcements and thus the divergence in returns. We therefore run security-level cross-sectional regressions of event returns on the predicted price impact π_i and a set of control variables

$$R_i^H = a_0 + a_1\,\pi_i + a_2\,u_i + a_3\,\log(cap_i) + a_4\,\beta_i^{mkt} + a_5\,\beta_i^F + a_6\,\text{Amihud}_i + \eta_i \quad (3.21)$$

For the purpose of these regressions, event returns are defined as the cumulative returns computed over the 10 trading days following the announcement ($H = 10$). We control for each security's weight in the purchase schedule of the BoJ (u), the natural logarithm of its market capitalization, its market beta, its Forex beta and its Amihud ratio as a proxy for illiquidity.

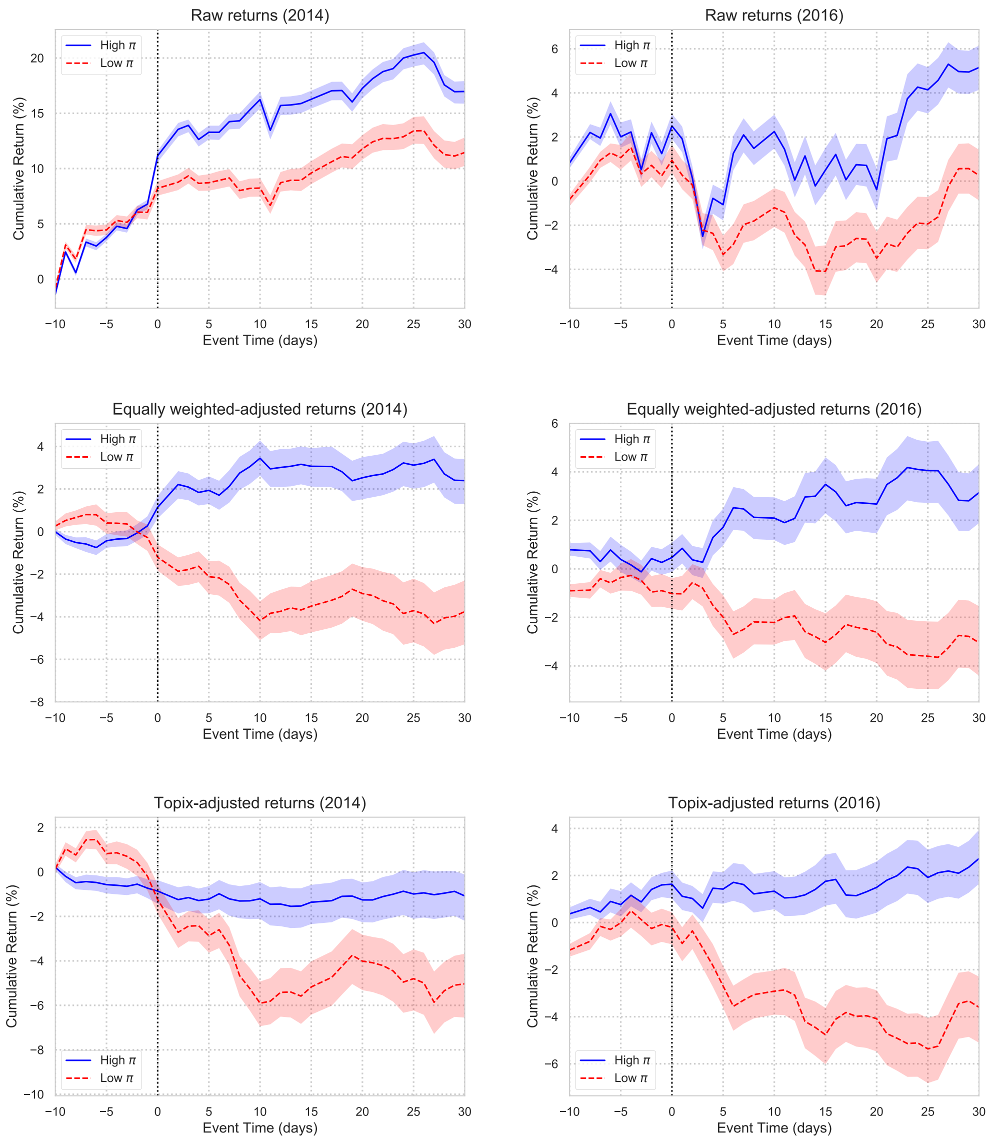

Figure 3.4: Cumulative returns of high versus low π stocks (in percentage). This figure shows the time series of the mean cumulative returns around the BoJ announcements of stocks with high predicted price impact π against that of low π stocks. The plots on the left refer to the announcement on October 31st, 2014, while those on the right show the reaction to the announcement on July 29th, 2016. The two top panels plot the unadjusted returns. In the four remaining panels returns are adjusted using a market model estimated in a window of one year, as described in Section 3.5.3. An equally-weighted portfolio of stocks in the TOPIX universe is used a proxy for the market portfolio in the middle panels, while the return of the TOPIX index is used in the bottom panels. The blue line is the average for the first quartile of the distribution (firms with the highest predicted price impact), while the red dashed line corresponds to the average for the last quartile (firms with the lowest predicted price impact). Bands represent bootstrapped 95% confidence intervals.

Panel A: October 31st, 2014

	Raw Returns				Abnormal Returns			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
π	57.86***	59.15***	31.92***	23.27***	36.06***	37.75***	39.95***	30.60***
	(8.59)	(8.00)	(4.40)	(3.49)	(4.94)	(4.81)	(5.54)	(4.62)
u		-0.00	-0.02***	-0.02***		-0.00	-0.02***	-0.02***
		(-0.95)	(-3.83)	(-3.43)		(-1.34)	(-3.62)	(-3.20)
Market Beta			0.040	0.025			-0.05	-0.07*
			(0.75)	(0.51)			(-1.48)	(-1.97)
Forex Beta			0.040*	0.043**			0.040*	0.041**
			(1.87)	(2.32)			(1.96)	(2.32)
log(Market Cap)			0.007	0.007			0.005	0.006
			(1.17)	(1.36)			(0.97)	(1.16)
Amihud			0.000	0.000			0.000	4.618
			(0.53)	(0.49)			(0.22)	(0.03)
Observations	1,851	1,851	1,807	1,701	1,851	1,851	1,807	1,701
R-squared	0.106	0.106	0.162	0.203	0.046	0.047	0.108	0.160
Industry FE	NO	NO	NO	YES	NO	NO	NO	YES

Panel B: July 29th, 2016

	Raw Returns				Abnormal Returns			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
π	14.07*	14.30*	12.17*	11.88	15.33*	16.33*	17.49**	16.69**
	(2.09)	(1.93)	(1.68)	(1.78)	(2.10)	(2.08)	(2.43)	(2.52)
u		-0.00	0.001	0.002		-0.00	0.004	0.004
		(-0.29)	(0.19)	(0.38)		(-1.33)	(0.56)	(0.67)
Market Beta			0.006	-0.00			0.002	-0.01
			(0.12)	(-0.06)			(0.06)	(-0.26)
Forex Beta			0.016	0.012			0.019	0.014
			(0.78)	(0.66)			(0.94)	(0.80)
log(Market Cap)			-0.00	-0.00			-0.00	-0.00
			(-0.10)	(-0.05)			(-0.58)	(-0.49)
Amihud			0.000	0.000			0.000	0.000
			(0.43)	(0.18)			(0.41)	(0.11)
Observations	1,905	1,905	1,839	1,734	1,905	1,905	1,839	1,734
R-squared	0.017	0.017	0.021	0.043	0.019	0.019	0.028	0.050
Industry FE	NO	NO	NO	YES	NO	NO	NO	YES

Table 3.2: Cross-sectional regressions. The tables report the regression coefficients of the cross-sectional regression of returns (in percentage points) on the predicted price impact π and a set of control variables (standardized). Regressions are run separately for the two events. The dependent variable in columns 1-3 is the cumulative raw return, while in columns 4-6 is the cumulative abnormal return with respect to the market model estimated in the pre-event window. Cumulative returns are computed over a 10 days horizon after the announcement date. t-statistics from placebo regressions are in parenthesis; asterisks denote conventional significance levels (***=1%, **=5%, *=10%) based on empirical p-values.

The Internet Appendix A reports summary statistics of the control variables by quartile of π. Consistent with the fact that the policy is heavily skewed towards Nikkei companies, which are on average larger than non-Nikkei ones, we find a positive correlation between π and market capitalization. Market capitalization is therefore an omitted variable in a regression of stock returns on π and we need to control for it. Moreover, the policy announcement could affect equity prices through its impact on the foreign exchange market. Since π is weakly negatively correlated with the Forex exposure β^F, if the yen depreciated as a consequence of the announcement, we would spuriously observe returns proportional to π. We therefore control for the exposure to the exchange-rate by adding β^F to the regressions. We also control for market betas. Notice though that there is no obvious relationship between π and market betas. We include the weights of the BoJ purchase schedule u_i to control for alternative explanations based on the direct effect of purchases in which LSAPs affect asset prices proportional to the amount purchased. Finally, we include industry fixed effects to some specifications, to make sure that our results hold within industries.

We run these regressions on the entire universe of TOPIX firms. Panel A of Table 3.2 investigates the cross-sectional effect of the BoJ announcement on October 31, 2014 (when the target purchase amount of ETFs was tripled), while Panel B analyzes event returns following the announcement on July 29, 2016 (when the target was further doubled). In either cases, no change was made to the weighting scheme of the purchases. In the first four columns the dependent variable is the cumulative raw return of the stock, while in the last columns the left-hand side variable is the cumulative abnormal return from a market model calculated using pre-event market betas.

On a given day, stock returns are expected to be correlated in the cross-section and therefore the OLS assumption of iid residuals is likely to be violated. We therefore run placebo regressions on the period from January 2009 to March 2013 to get the empirical distribution of the coefficients in the absence of policy shocks, which we use to compute robust standard errors. The placebo event days are chosen randomly on non-overlapping periods to ensure that the empirical distribution is constructed from independent draws. For regressions involving short-horizon returns (up to 3 months) we impose that placebo event periods do not include BoJ meetings on which important monetary policy announcements were made. Namely, we exclude the meetings of February 1st 2013, March 25th 2013, June 18th 2012 and the announcement of the post-tsunami intervention in March 14th 2011. On all regression tables of this paper we report the empirical p-values computed using this methodology.

Consistent with Proposition 1, the coefficient on the predicted price impact π is positive and significant across specifications and events. For the 2014 announcement, the baseline specification with raw returns reported in columns (1) shows a remarkable R^2 above 10%, suggesting that our expected price impact π is crucial to explain the heterogeneity of event returns. As it was already visible from the plots in the previous section, the results are weaker for the 2016 event. In particular, the portion of explained variance for the 2016 policy announcement is significantly lower, consistent with the smaller change in the target purchase amount. Still, the coefficient on π is positive and significant at the 10% confidence level in most specifications. The coefficient turns however insignificant when we include industry fixed effects in the specification with raw returns.

Results show that the effect of π is robust to the inclusion of the vector u of purchased amounts. This horse race provides additional support for the portfolio-balance channel against a local channel where spillovers are negligible. The model predicts that the effect of u should be insignificant once we control for π. This is indeed what we find in the second specification of each panel. The coefficient on u turns however negative and significant in columns (3) and (4) of panel A. We find that this is due to the inclusion of the control for market capitalization since u and market cap are highly correlated, as it is natural to expect. Still, this does not affect the size and the significance of the coefficient on π.

Results also show that controlling for the exposure to the exchange rate does not impair the significance of the coefficient on π. The coefficient on β_F is positive and significant in 2014, when the BoJ announcement was followed by a rise in the Forex. In 2016, on the other hand, the coefficient on β_F is not significant, consistent with the fact that the Forex did not move significantly (see Figure A.5 in the Appendix).

In column (2) of the regression using raw returns as dependent variable, the coefficient on π drops significantly. This is, as expected, due to the fact that control variables play an important role in explaining cross-sectional returns variation, as documented by a significantly larger coefficient of determination. In particular, market beta, Forex beta and market capitalization incrementally increase the regression's R^2 and dampen the coefficient on π. In columns (3) the number of observations drops slightly because of missing data on trading volume needed to estimate the Amihud ratio. In columns (4) the sample is further reduced because of missing information on industry classification.

3.6.3 Time-Series Pattern

We now turn to the long-run predictions of the model summarized in Proposition 2. We estimate the cross-sectional model specified in equation (3.21) at different horizons H over which cumulative returns are calculated. Results are reported in Table 3.3. At portfolio level, Figure 3.4 suggests that the cross-sectional effect of the BoJ announcements is long-lasting and weakly increasing over time. The regression analysis confirms that the evidence holds at stock-level and after controlling for security specific characteristics. The vector π is positively and significantly related to cross-sectional stock returns at every horizon H after the announcement. In other words, the model implied changes in systematic risk estimated ex-ante are a significant predictor of post-event abnormal returns across stocks.

We find no evidence of reversal of the initial price impact even one year after the event. The absence of reversal is a key prediction of the portfolio-balance channel. Since the shock to supply induced by the BoJ has a unique cross-sectional shape, it is unlikely that the observed effect is due to shocks other than the purchase program, suggesting a causal effect of the policy. Still, we cannot completely rule out other policy transmission mechanisms. In Section 3.7, we address the possibility that (part of) the effect might be due to continuous price pressure that prevents prices from reverting to the pre-announcement level.

	Abnormal Returns 2014						Abnormal Returns 2016					
	5	10	21	63	126	252	5	10	21	63	126	252
π	17.78**	39.95***	28.31***	72.21***	171.9***	305.5***	17.43**	17.49**	33.70***	40.86**	93.52***	120.2***
	(3.09)	(5.54)	(2.89)	(4.52)	(8.97)	(11.91)	(3.03)	(2.43)	(3.44)	(2.56)	(4.88)	(4.69)
u	-0.00	-0.02***	-0.01	-0.03***	-0.03**	-0.07***	0.011**	0.004	0.026**	0.016	0.052***	0.118***
	(-0.30)	(-3.62)	(-1.21)	(-3.20)	(-2.46)	(-3.43)	(2.22)	(0.56)	(2.76)	(1.48)	(3.31)	(5.13)
Market Beta	-0.02	-0.05	-0.07	-0.16*	-0.29**	-0.40***	0.026	0.002	0.025	0.025	0.041	0.068
	(-0.84)	(-1.48)	(-1.36)	(-2.02)	(-2.04)	(-2.51)	(0.86)	(0.06)	(0.47)	(0.32)	(0.28)	(0.43)
Forex Beta	0.036**	0.040*	0.101***	0.097**	0.043	-0.13	-0.00	0.019	0.061*	0.070*	0.220***	0.216
	(2.17)	(1.96)	(3.83)	(2.29)	(0.61)	(-0.02)	(-0.23)	(0.94)	(2.30)	(1.67)	(3.10)	(0.04)
log(Market Cap)	0.001	0.005	0.001	0.001	0.005	-0.00	-0.00*	-0.00	-0.01*	-0.02	-0.07***	-0.11***
	(0.32)	(0.97)	(0.17)	(0.06)	(0.25)	(-0.27)	(-1.60)	(-0.58)	(-1.81)	(-1.80)	(-3.52)	(-6.82)
Amihud	0.001	0.000	-1.02	0.001	0.016***	0.018***	0.000	0.000	-0.00	-0.00	-0.00**	-0.00
	(0.78)	(0.22)	(-0.01)	(0.70)	(6.27)	(7.17)	(0.47)	(0.41)	(-0.47)	(-0.92)	(-3.12)	(-2.16)
Observations	1,807	1,807	1,807	1,807	1,807	1,807	1,839	1,839	1,839	1,839	1,839	1,839
R-squared	0.055	0.108	0.073	0.098	0.153	0.119	0.051	0.028	0.079	0.077	0.178	0.140

Table 3.3: Cross-sectional regressions over different horizons. The table report the coefficients of cross-sectional regressions of cumulative returns (in percentage points) computed at different horizons on the predicted price impact π and a set of control variables (standardized). Regressions are run separately for the two events. The dependent variable is the cumulative abnormal return with respect to the market model estimated in the pre-event window. t-statistics from placebo regressions are in parenthesis; asterisks denote conventional significance levels (***=1%, **=5%, *=10%) based on empirical p-values.

A second finding from Table 3.3 is that the estimated coefficients on π are generally increasing in H.[10] Post-event returns in the same direction of the announcement effects are predicted by the model through the decrease in residual duration of the program. However, as we discuss in Section 3.4, this effect is expected to be small for realistic values of the interest rate. The model produces a similar post-event return pattern when beliefs about the size of the program (λ_t) are increasing over time. This suggests that investors might be extrapolating current purchases above and beyond the policy horizon or that they might not believe to a full commitment of the central bank to the announced purchase target at first, but slowly update their beliefs. In the current setting we cannot disentangle between these explanations, nor convincingly claim that post-event returns are in fact driven by updating in beliefs. The question is therefore open for future research.

In the Internet Appendix we present some robustness evidence. We show that the observed price impact cannot be explained by industry effects, nor it is simply driven by an over-performance of Nikkei stocks versus non-Nikkei stocks. The results remain largely unchanged across specifications.

3.6.4 Quantification of Portfolio-Balance Effects

In this section we propose a simple back-of-the-envelope calculation to try to quantify the net aggregate portfolio balance effect of the BoJ intervention from the coefficient estimated in the cross-section. From this quantity we then derive an estimate of the aggregate elasticity of equity demand curves.

For this calculation we want to use the average effect of the policy across the two events, so we first run again our main regression in equation (3.21) over the pooled sample, including event fixed-effects FE_e to allow for a different intercept across the two announcements. Precisely, we estimate the following regression model for each daily horizon $h \in 1, \ldots, 252$

$$R^h_{i,e} = \beta^h \pi_{i,e} + \gamma^h X_{i,e} + \delta^h \, \text{FE}_e + \varepsilon_{i,e} \tag{3.22}$$

where X is a vector of control variables that depends on the regression specification and $e \in (2014, 2016)$ is an index numbering the events. Since we are considering both events together, we need to rescale the π vectors to take into account the different magnitude of the announcements. Therefore we multiply π_{2014} by 3 and π_{2016} by 6 to reflect the magnitude of the target amount announced by the BoJ in the two events, respectively. We include market capitalization in each specification to control for the size factor, which is expected to become more relevant as the horizon increases. In the second specification we also control for market and Forex betas. In

[10]Notice that the cumulative returns on the left-hand side of the regression are computed as cumulative sums rather than cumulative products in order to avoid a mechanical effect when increasing the horizon.

	1 week	1 month	3 months	6 months	1 year
(1) Baseline	3.54	10.10	10.28	25.17	22.32
(2) Control for market and Forex	2.23	7.72	9.53	22.08	22.45
(3) Control for market, Forex and liquidity	1.56	7.17	9.56	19.80	22.05

Table 3.4: Portfolio Balance Effects. The table presents the estimated net portfolio balance effect on the market, expressed in basis points per Trillion Yen invested by the central bank into the ETF purchase program. We report point estimates for the net effect impounded into prices over increasing horizons, from three models employing different sets of control variables defined in the text.

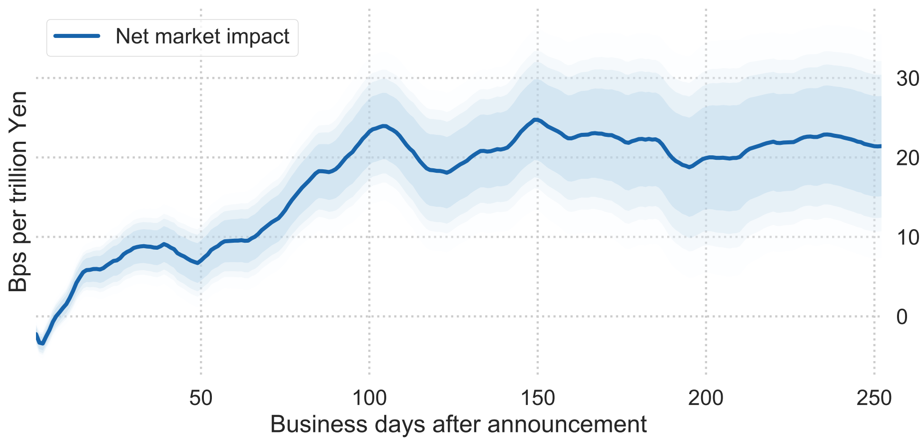

Figure 3.5: Portfolio Balance Effects. This figure plots the time-series evolution of the estimated portfolio balance effect induced by the BoJ purchase program, expressed in basis points per trillion Yen invested. The estimates are based on specification (3), which includes controls for stocks liquidity, market beta and exposure to the US-JPN Forex exchange rate. Thus the estimated market impact can be interpreted as the counter-factual policy effect, net of of alternative channels and confounding factors. Shaded areas denote 10%, 5% and 1% confidence intervals.

the third specification we additionally include each stock's Amihud ratio to control for liquidity.

Given $\hat{\beta}^h$ from the estimation, the predicted net return through the portfolio balance channel for security i is $\hat{R}^h_{i,e} = \hat{\beta}^h \pi_{i,e}$ [11]. To aggregate the effect at market level, we calculate for each event $e \in (2014, 2016)$ the predicted market return as the value-weighted sum of security level predicted returns at every horizon

$$\hat{R}^h_e = \hat{\beta}^h \sum_i w_{i,e} \pi_{i,e} \tag{3.23}$$

We then divide by the capital commitment by the central bank to obtain the induced market return per trillion yen. Considering the two-year policy horizon, this

[11] Notice that the estimated $\hat{\beta}^h$ allows us in principle to compare the impact of alternative purchase schedules u' that the central bank could have implemented, conditional on the same covariance matrix Σ. In this section we are interested in the estimated portfolio balance effect of the actual purchase portfolio.

amounts to 6 trillion Yen for 2014 and 12 trillion Yen for 2016, with the underlying assumptions that each announcement was completely unexpected and that investors are reacting to the announced program size. Thus, the per yen estimated average market return induced by the policy through the portfolio-balance channel is calculated as

$$\hat{R}^h = \frac{1}{2}\left(\hat{R}^h_{2014}/6 + \hat{R}^h_{2016}/12\right) \tag{3.24}$$

Results of this exercise are reported in Table 3.4 for the three specifications. The last column shows an estimated long-term impact of about 22 basis points increase in market value per trillion yen employed. With about ¥500 trillion of total market capitalization, this implies an elasticity close to one since each yen invested translates into an increase of the market valuation by roughly one yen.

Figure 3.5 plots the time-series evolution of the point estimate for the third specification, showing that the portfolio balance effects are slowly impounded into prices. Consistent with the qualitative prediction of our model, a momentum-like pattern is visible over the first 100 trading days following the announcement.

Notice that the quantity we are estimating in this section is the aggregate portfolio-balance effect of the policy on the returns of stocks that are included in the TOPIX index. While the policy is expected to have additional effects through different channels, our empirical methodology allows us to identify and quantify the portfolio balance channel. In turn, this allows us to derive an estimate of the price elasticity of the demand curve for stocks, which has to be intended as local to the First Section of the TSE. We acknowledge that the policy might have produced spillover to other asset classes (this is indeed a prediction of the model), but we are not considering them in this paper.[12]

The derivation of the aggregate portfolio balance effect described in this section relies on two main assumptions. First, it depends on our assumption that the representative agent in our model re-invests the proceeds from the sale of stocks to the central bank at the constant risk-free rate.[13] The key point here, is that the risk-free asset is uncorrelated with the Japanese stocks. While this is not a concern for the identification of the portfolio-balance effect, it is more problematic when we try to quantify the effect, since our approach might be providing a biased estimate of the aggregate effect if the assumption is not valid.

To understand the direction and magnitude of the potential bias, in the Internet Appendix we extend the model to allow the representative agent to re-invest the proceed in a security correlated with the targeted assets. We derive an expression of the resulting bias if we incorrectly assume the above mentioned assumption, which can be seen as a form of market segmentation. The bias is a function of the market

[12] Spillovers to unaffected stocks is already a key point in Greenwood, 2005 and is to be expected in this setting as well. We believe that the cross-sectional heterogeneity among the stocks in the TOPIX is sufficient to support our arguments. Moreover, the TOPIX index covers all First Section companies in the Tokyo Stock Exchange (TSE), which are the majority of public companies in Japan and is by far largest section of the TSE in terms of market capitalization and trading volume.

[13] We thank the Anonymous Referee for pointing out this issue.

weighted average of the covariance of the omitted variable and π, where the omitted variable is the vector of covariances between the re-investment security and the stocks. The sign of the bias is ambiguous and depends on Σ, u, the market weights and the re-investment security. We therefore run simulations of the model using the parameters estimated in the data and assuming different re-investment securities, namely S&P500, 10-year US Treasury bonds and 10-year JGBs. The estimated bias is positive using the S&P500 and negative using long-term government bonds. The magnitude of the bias is relatively small, around (positive or negative) 10%. Depending on which direction the bias is going, the estimated elasticity of 1 might be slightly over- or under-estimating the true elasticity of Japanese equities.

A second reason why our approach might be delivering a biased estimate of the aggregate effect is that in our calculation we are assuming that the market is reacting to the announced size of the program. If the market is in fact reacting to expectations of a smaller program, either because investors think the BoJ will not reach the announced target or because it will soon unwind its portfolio, then the estimated elasticity of 1 represents an upper-bound for the true price elasticity. Vice versa, the estimated elasticity of 1 would be a lower-bound if investors believe that the BoJ will continue the purchase program beyond M.

3.7 Portfolio Rebalancing or Price Pressure?

The previous section shows that the reaction of stock prices to the upward revisions of the purchase target is consistent with a portfolio balance channel. Even at long horizons, the expected change in systematic risk is key to explain the cross-sectional variation in returns after the policy announcement. We interpret the persistence of the effect as evidence of downward sloping long-run demand curves.

An alternative explanation for the observed persistence relies on the continued pressure exercised by the BoJ through repeated purchases. If short-run demand curves are downward sloping, abnormal volumes induced by the BoJ during intervention days might push prices above fundamentals. Such effects are usually motivated by limits-to-arbitrage and are expected to revert quickly. However, as the central bank is expected to buy repeatedly, arbitrageurs may refrain from betting against mispricings and fail to bring prices back to fundamentals. If this was the case, the absence of reversal could not be interpreted as evidence for long-run demand curves sloping down. In the spirit of D'Amico and King, 2013, we will refer to this kind of dynamics as *flow effect* of the policy. This naming highlights that under this explanation the price impact is caused directly by the trading volume (or *flow*) of the central bank, rather than by reduction in systematic risk which underlies the portfolio balance channel.

The repeated price pressure story implies that we should observe higher positive abnormal returns on intervention days and that these should be proportional in the cross-section to the abnormal trading volume generated by the intervention. In this section we first introduce a reduced form model that exploits the time-series and cross-sectional variation in daily purchases by the BoJ to estimate the flow effect of the policy. We then use the predicted returns from that model to remove the flow effect component from stock returns. Finally, we re-run the analysis of Section 3.6.2 on these *net* returns. By comparing the coefficient estimated in this way to the one in the previous section, we can assess how much of the observed price impact and its persistence is due to repeated price pressure rather than the portfolio balance mechanism.

Evaluating the relative magnitude of these two channels is essential to draw conclusions on the elasticity of long-run demand curves for stocks. The distinction between the two explanations has also practical consequences for policy makers regarding the exit strategy from the purchase program. A repeated price pressure story predicts prices to revert as soon as the buying pressure from the central bank stops, making the accumulated size of the balance sheet *de facto* irrelevant beyond that point. On the contrary, in the model of Section 3.4 the aggregate impact of the policy is unaffected by the timing of the purchases. In the extreme case where the central bank buys everything on the announcement day, the model predicts an immediate, complete and permanent price adjustment.

3.7.1 Purchase Frequency and Volumes

The QQE was announced on April 4, 2013, and the asset purchases were then gradually carried out. In its official statements, the BoJ does not commits itself to any particular purchase frequency and does not reveal in advance the days in which it will buy. Ex-post, we can see from Panel A of Figure 3.6 that the bank has been buying fairly consistently once to twice a week over the sample period. The blue crosses in Panel B of Figure 3.6 indicate intervention days. On the y-axis we report the ratio between the amount purchased and the aggregate trading volume in the underlying stock market on that day.

Since the purchase frequency remained stable over the policy horizon, the upward revisions of the annual target in October 2014 and in July 2016 translated into an increase of the daily purchased amount. However, even in the last period, the quantity purchased by the BoJ represented less than 5% of the daily market volume, a threshold which is often used by practitioners as guideline for when a trade is expected to have a significant price impact. At the stock level, the purchases by the BoJ account for more than 5% of the daily trading volume on average only for 5.2% of the targeted stocks.

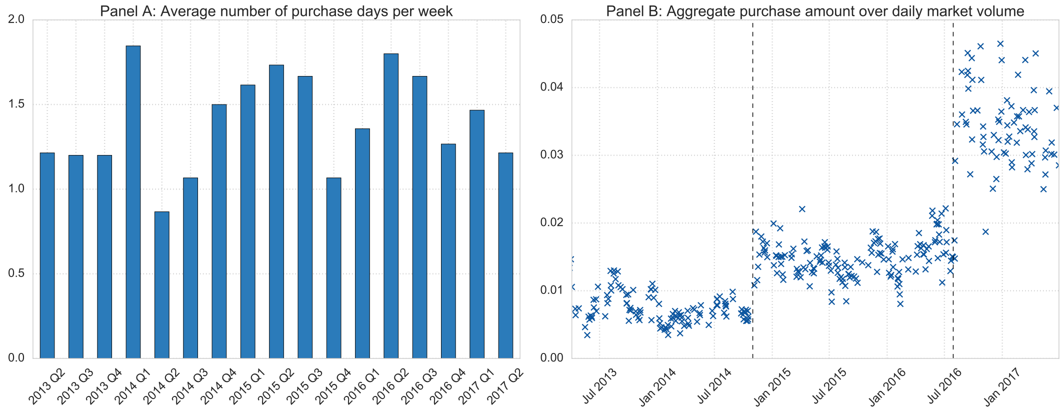

Figure 3.6: Purchase Frequency and Volume. Panel A plots the average number of purchase days per week at quarterly frequency. Panel B plots the ratio between the yen amount purchased by the BoJ on a given day and the aggregate trading volume in yen on that day. The aggregate trading volume is computed as the sum of the trading volume of the securities targeted by the policy.

3.7.2 Empirical Setup and Results

To quantify the direct price impact of purchases we estimate a dynamic model in the spirit of Eser and Schwaab (2016) that relates daily stock returns to daily flows from the central bank. The model is specified as

$$AR_{i,t} = \alpha + \beta_0 AV_{i,t} + \beta_1 AV_{i,t-1} + \beta_2 \left(\sum_{k=2}^{K} \rho^{k-2} AV_{i,t-k} \right) + \varepsilon_{i,t} \tag{3.25}$$

where the left-hand side variable $AR_{i,t}$ is the daily abnormal return of stock i relative to the market model, estimated following the methodology outlined in Section 3.5.3. On the right-hand side, the BoJ-induced abnormal volume $AV_{i,t}$ is defined by

$$AV_{i,t} := \frac{\text{BoJ Flow}_{i,t}}{\text{E}\left[\text{Volume}_{i,t}\right]} \tag{3.26}$$

and measures the size of the purchased amount of stock i on day t relative to the average market volume of that stock. The purchased amount BoJ Flow$_{i,t}$ is computed as $\frac{1}{2}(w_{i,T} + w_{i,N})A_t$, where $w_{i,T}$ is the weight of stock i in the TOPIX index, $w_{i,N}$ is the weight of stock i in the Nikkei 225 index and A_t is the value of ETFs purchased by the BoJ on day t. Here we assume that each trade of the BoJ in the ETF market translates into proportional shocks to the underlying basket on the same day.[14]

[14]Underlying securities inherit shocks that occur in the ETF market both through primary market arbitrage as well as through the arbitrage that takes place continuously in the secondary market and that is carried out by hedge funds and high-frequency traders (Ben-David, Franzoni, and Moussawi, 2018). Secondary market arbitrageurs make profits by opening their positions when the price of the ETF deviates from NAV and holding them until prices converge. Our identification of the flow effect of the Policy relies on arbitrageurs trading on the same day as the BoJ. For secondary market arbitrage, this is a reasonable assumption. Competition among arbitrageurs implies that hedge funds and high-frequency traders will open their positions as soon as they observe the ETF trading at a premium over the NAV. That such arbitrage opportunities exist on the days when the BoJ buys is consistent with the

The average daily volume $\mathrm{E}\left[\mathrm{Volume}_{i,t}\right]$ is estimated over a backward-looking window of six months excluding days in which the BoJ is intervening. On non-purchase days the abnormal volume is therefore zero for every stock in our sample, while it is strictly positive on purchase days.

The model includes lagged values of AV to capture the permanent component of the price pressure, net of transitory and delayed effects of purchases. The long-run effect of the flow-induced price impact can be computed from the estimated coefficients as

$$F = \beta_0 + \beta_1 + \beta_2 \left(\sum_{k=2}^{K} \rho^{k-2} \right) \tag{3.27}$$

The parameter $\rho \in (0, 1)$ determines how long it takes for prices to adjust following an intervention. If ρ is close to zero the dynamic of the flow effect is exhausted after two days. $F \approx 0$ implies that temporary price impacts, if any, are fully reverted. This in turn would mean that the price pressure story does not contribute to explain the persistence of the policy impact documented in Section 3.5. On the contrary, $F > 0$ implies that (part of) the persistence attributed to the portfolio balance mechanism might be due to the direct impact of the flow of BoJ purchases.

The identification of the direct impact of the purchases (flow effect) in this panel regression framework relies both on the exogeneity of the cross-sectional variation of the purchases and on the predetermination of the purchase amounts with respect to prices. The exogeneity in the cross-section is discussed extensively in Section 3.2 and mainly relies on the fact that the weighting system of the Nikkei 225 introduces significant variation in the cross-section of purchases that is unrelated to firms' fundamentals. Predetermination of the purchases is not straightforward in the current context. The criteria used by the BoJ to decide whether and how strongly to intervene on a particular day are not public information, however there are reasons to believe that the BoJ tends to intervene on days when the market is falling. In fact, the median stock return is significantly lower on intervention days (-0.6%) relative to non-intervention days (0.3%). To tackle this potential endogeneity of BoJ flows we specify the regression model in terms of abnormal returns. Given that the BoJ might be using the return on the market as a signal for whether to intervene, removing the contemporaneous return on the market should mitigate the issue. The fact that mean and median abnormal returns are not significantly different from zero in both intervention and non-intervention days supports our claim.

We do not include an announcement dummy in the specification because on those days no ETF purchases were made by the BoJ. Looking at the time series of BoJ purchases, we see that the bank intervened two weeks before and one week after the first upward revision of the purchase target on October 31, 2014. Similarly, no purchases were made on July 29, 2016. Purchases are registered on the previous day and four days after.

evidence in Figure 3.3, since growth in AUM is consistent with upward pressure on ETF prices. The results reported in Table 3.5 provide further support for the validity of this assumption since they show that most of the price impact seem to take place on the event day and the day after.

		Panel A					Panel B		
Model	K	β_0	β_1	β_2	ρ	F	$\tilde{a}_1$	Flow Effect	Port Balance
(1)	0	0.011				0.011	138.644	5.63%	94.37%
		(9.356)					(11.175)		
(2)	1	0.004	0.015			0.019	132.288	9.96%	90.04%
		(3.314)	(11.358)				(10.662)		
(3)	2	0.004	0.015	-0.001		0.018	132.608	9.74%	90.26%
		(3.358)	(10.979)	(-0.679)			(10.688)		
(4)	5	0.004	0.015	-0.001	0.001	0.018	132.567	9.77%	90.23%
		(3.396)	(10.997)	(-0.694)	(0.035)		(10.685)		
(5)	10	0.004	0.015	-0.001	0.001	0.018	132.567	9.77%	90.23%
		(3.396)	(10.997)	(-0.694)	(0.039)		(10.685)		

Table 3.5: Flow Effect The table reports results from the estimation of the dynamic model described in (3.25), where a different value for the number of lags K is used in each specification. The models are estimated with maximum likelihood assuming normally distributed error terms and constraining the persistence parameter ρ in the unit interval. Panel A presents the estimated model parameters and the implied long-run effect F. Panel B shows OLS estimates of the coefficient $\tilde{a}_1$ resulting from a cross-sectional regression of cumulative abnormal returns, purified from the estimated flow effects, on the predicted price impact π resulting from the portfolio balance model of Section 3.4. The decomposition into *flow* and *portfolio balance* components is obtained by comparing $\tilde{a}_1$ with the coefficient a_1 from Section 3.6.2 based on standard CARs.

We estimate five different specifications of model (3.25). We start considering only contemporaneous volumes ($K = 0$), then we augment the specification to K equal to $1, 2, 5$ or 10. The estimated parameters are reported in Panel A of Table 3.5 together with the implied long-run impact. The positive coefficients on β_0 and β_1 suggest that abnormal returns are significantly higher during purchase days for stocks experiencing a higher degree of buying pressure. The negative but not significant value of β_2 and a persistence parameter ρ close to zero suggest that such a price impact is not reverted in the next trading weeks and give rise to a positive long-run component F in every specification.

The results indicate a positive and persistent flow effect of the policy, which might lead to an overestimation of the portfolio balance channel in the previous section. To quantify the consequences of not taking flows into account, we construct the flow induced returns as the fitted values of the estimated model

$$\widehat{AR}_{i,t}^{Flow} = \hat{\beta}_0 AV_{i,t} + \hat{\beta}_1 AV_{i,t-1} + \hat{\beta}_2 \left(\sum_{k=2}^{K} \hat{\rho}^{k-2} AV_{i,t-k} \right) \tag{3.28}$$

which we subtract from stock returns to remove the direct impact of the BoJ purchases

$$\widetilde{AR}_{i,t} = AR_{i,t} - \widehat{AR}_{i,t}^{Flow} \tag{3.29}$$

We then estimate our main regression (3.21) using $\widetilde{AR}$ instead of AR, computing the cumulative abnormal returns over a one-year horizon following the two event dates and we regress them on the predicted price impact vector π. We pool the

2014 and 2016 events together to obtain a unique estimate $\tilde{a}_1$ of the price impact of the policy through the portfolio balance channel. The ratio between $\tilde{a}_1$ and its counterpart $\hat{a}_1$ obtained estimating the model with the cumulative returns computed from AR, gives us the fraction of the estimated portfolio balance impact that might be explained by the price pressure channel.

Panel B of Table 3.5 summarizes the results of this second step, showing that the fraction of the observed cross-sectional pattern explained by the price pressure channel ranges between 5% and 10% depending on the specification. It must be noted that these figures represent upper bounds for the persistent flow effect of the policy, since this might be amplified by expectation updates consistent with the portfolio balance model, if investors learn about the commitment of the central bank through the realization of its purchases.

Taken together, the results of this section suggest that the price pressure generated by the central bank at the stock level plays a limited role in explaining the impact of the policy. We conclude that the observed cross-sectional pattern of stock returns is mostly generated by the portfolio balance channel rather than continued price-pressure arising from the central bank flows.

3.8 Policy Implications

In this section we discuss the policy implications of our results. Based on our theoretical framework, we show formally that the heterogeneity uncovered by our empirical analysis could be avoided if the central bank would buy the value-weighted market portfolio, since this would lead to a homogeneous reduction of firms' cost of capital in the cross-section.

Recall that in our model the cost of capital of each firm is proportional to its marginal risk contribution to the market portfolio (systematic risk). Formally, the vector of risk premia prior to the BoJ intervention is proportional to VQ, where V is the variance-covariance matrix of fundamentals and $Q \in \mathbb{R}^n$ is the vector of shares outstanding.

As soon as the central bank purchases a quantity $q \in \mathbb{R}^n$, the cost of capital is affected and converges to $V(Q - Mq)$. In particular, firm i experiences a percentage shift in its perceived cost of capital equal to

$$\Delta k_i = \frac{(V(Q - Mq))_i}{(VQ)_i} - 1 \tag{3.30}$$

Notice that Δk_i is not necessarily negative, thus some firms may experience an increase in their financing costs ($\Delta k_i > 0$), even if the central bank buys some of their shares ($q_i > 0$).

It follows that a homogeneous impact on risk premia can be achieved with a vector of purchases proportional to Q. If the purchase schedule is $q^* = aQ$ for $a \in \mathbb{R}$, the effect on firm i is

$$\Delta k_i^* = \frac{(V(Q - Mq^*))_i}{(VQ)_i} - 1 = \frac{((1 - Ma)VQ)_i}{(VQ)_i} - 1 = \frac{((1 - Ma)VQ)_i}{(VQ)_i} - 1 = -Ma \tag{3.31}$$

which does not depend on i and is thus homogeneous across companies.

In the case of Japan, a purchase schedule q parallel to Q corresponds to the BoJ limiting its purchases of ETFs to those tracking the value-weighted TOPIX Index[15]. Buying ETFs tracking the price-weighted Nikkei 225, on the other hand, introduces a component in q which is orthogonal to Q. This, in turn, leads to heterogeneous consequences for firms financing costs, which can be interpreted as a distortion of the market allocation mechanisms. Figure A.6 shows that the distortion is evident also at the industry-level.

Under the assumption that a homogeneous effect is the preferred outcome of the policy, we infer from the model that the central bank should stop buying Nikkei-indexed ETFs. More precisely, the central bank should schedule future purchases with the objective of re-shaping its equity portfolio in a value-weighted fashion.

A change of policy in this direction was solicited by a number of critics of the purchasing program, and on September 2016 the BoJ changed the guidelines for its asset purchases, reducing the share of capital flowing to ETFs tracking the Nikkei 225 Index and increasing its holdings of ETFs tracking the TOPIX. This brought the cross-sectional allocation of capital closer to what market capitalization would justify.

To date, the BoJ has not completely abandoned the price-weighted Nikkei Index, nor it is bringing its already accumulated holdings towards value-weighted proportions. According to our model, the BoJ should make sure to bring its holdings proportional to companies market capitalizations if it wants to amend the allocational side-effects of the policy.

3.9 Conclusion

In this paper we study asset pricing implications of the ETF purchase program undertaken by the BoJ since April 2013. The analysis is supported by a dynamic asset pricing model, featuring multiple assets with time-varying supply due to open market operations of the central bank.

To identify the net portfolio balance effect of the policy our empirical analysis exploits the exogeneity and the cross-sectional dimension of the BoJ's purchase schedule, which mitigates endogeneity problems characteristic of other studies.

[15] A purchase of $u' = aW_{\text{Topix}}$ in Yen corresponds to $u = aQ$ in shares, since the TOPIX is value-weighted.

We show that the intervention has a positive and persistent effect on domestic equity prices, thus reducing the cost of equity capital of domestic companies. We provide empirical evidence that the effect is consistent with a portfolio balance channel both in the cross-section and in the time-series.

This evidence suggests that demand curves for stocks are downward sloping in the long-run. We estimate an economically significant increase of 22 basis points in aggregate market valuation per trillion Yen invested into the program, which corresponds to a price elasticity of 1. The mechanism behind downward sloping demand curves in the model is through the change in the structure of systematic risk held by the private sector induced by the central bank's intervention. This change in the composition of risk leads to a new discount factor and consequently to price adjustments.

We also show that the outright purchases of the BoJ generate positive and persistent pressure on prices. Our estimates of the portfolio-balance channel remain significant after accounting for these flow effects of the policy.

Our results shed light on the side-effects of the LSAP, uncovering a highly heterogeneous impact in the cross-section of firms' cost of equity capital, both at the firm and at the industry level. Using our theoretical framework to evaluate the impact of arbitrary purchase schedules, we find that the observed heterogeneity in the price effects mainly arises from the weight given to the Nikkei 225 price-weighted index. Capital injections shaped according to market weights would instead induce a cross-sectionally homogeneous change in the cost of capital.

References

Bagwell, Laurie Simon (1992). "Dutch auction repurchases: An analysis of shareholder heterogeneity". In: *The Journal of Finance* 47.1, pp. 71–105.

Barberis, Nicholas and Richard Thaler (2003). "A survey of behavioral finance". In: *Handbook of the Economics of Finance* 1, pp. 1053–1128.

Ben-David, Itzhak, Francesco Franzoni, and Rabih Moussawi (2018). "Do ETFs increase volatility?" In: *The Journal of Finance* 73.6, pp. 2471–2535.

Bernanke, B. S. (2009). *Federal Reserve policies to ease credit and their implications for the Fed's balance sheet: a speech at the National Press Club Luncheon, National Press Club, Washington, DC, February 18, 2009*. 443.

BoJ (2013). "Introduction of the Quantitative and Qualitative Monetary Easing, April 4th". In:

Brunner, Karl and Allan H Meltzer (1973). "Mr. Hicks and the "Monetarists"". In: *Economica* 40.157, pp. 44–59.

Buraschi, Andrea and Paul Whelan (2015). "Bond Markets and Unconventional Monetary Policy". In: *Handbook of Fixed-Income Securities* First Edition.

Coval, Joshua and Erik Stafford (2007). "Asset fire sales (and purchases) in equity markets". In: *Journal of Financial Economics* 86.2, pp. 479–512.

Culbertson, John M (1957). "The term structure of interest rates". In: *The Quarterly Journal of Economics* 71.4, pp. 485–517.

Curdia, Vasco and Michael Woodford (2011). "The central-bank balance sheet as an instrument of monetarypolicy". In: *Journal of Monetary Economics* 58.1, pp. 54–79.

D'Amico, Stefania and Thomas B King (2013). "Flow and stock effects of large-scale treasury purchases: Evidence on the importance of local supply". In: *Journal of Financial Economics* 108.2, pp. 425–448.

Eggertsson, Gauti B and Michael Woodford (2004). *Optimal monetary policy in a liquidity trap*. National Bureau of Economic Research (No. w10840).

Eser, Fabian and Bernd Schwaab (2016). "Evaluating the impact of unconventional monetary policy measures: Empirical evidence from the ECBs Securities Markets Programme". In: *Journal of Financial Economics* 119.1, pp. 147–167.

Frankel, Jeffrey A (1985). "Portfolio crowding-out, empirically estimated". In: *The Quarterly Journal of Economics* 100.Supplement, pp. 1041–1065.

Gagnon, Joseph et al. (2010). *Large-scale asset purchases by the Federal Reserve: did they work?* FRB of New York Staff Report 441.

Greenwood, Robin (2005). "Short-and long-term demand curves for stocks: theory and evidence on the dynamics of arbitrage". In: *Journal of Financial Economics* 75.3, pp. 607–649.

— (2007). "Excess comovement of stock returns: Evidence from cross-sectional variation in Nikkei 225 weights". In: *The Review of Financial Studies* 21.3, pp. 1153–1186.

Hamilton, James D and Jing Cynthia Wu (2012). "The effectiveness of alternative monetary policy tools in a zero lower bound environment". In: *Journal of Money, Credit and Banking* 44.s1, pp. 3–46.

Harris, Lawrence and Eitan Gurel (1986). "Price and volume effects associated with changes in the S&P 500 list: New evidence for the existence of price pressures". In: *the Journal of Finance* 41.4, pp. 815–829.

Hau, Harald, Massimo Massa, and Joel Peress (2009). "Do demand curves for currencies slope down? Evidence from the MSCI global index change". In: *The Review of Financial Studies* 23.4, pp. 1681–1717.

He, Zhiguo and Arvind Krishnamurthy (2013). "Intermediary asset pricing". In: *The American Economic Review* 103.2, pp. 732–770.

Joyce, Michael et al. (2011). "The financial market impact of quantitative easing in the United Kingdom". In: *International Journal of Central Banking* 7.3, pp. 113–161.

Joyce, Michael et al. (2012). "Quantitative easing and unconventional monetary policy – an introduction". In: *The Economic Journal* 122.564.

Kaul, Aditya, Vikas Mehrotra, and Randall Morck (2000). "Demand curves for stocks do slope down: New evidence from an index weights adjustment". In: *The Journal of Finance* 55.2, pp. 893–912.

Krishnamurthy, Arvind and Annette Vissing-Jorgensen (2011). *The effects of quantitative easing on interest rates: channels and implications for policy*. National Bureau of Economic Research.

— (2013). "The ins and outs of LSAPs". In: *Federal Reserve Bank of Kansas City's Jackson Hole Symposium on the Global Dimensions of Unconventional Monetary Policy, Jackson Hole, Wyoming, August*.

Ledoit, Olivier and Michael Wolf (2004). "A well-conditioned estimator for large-dimensional covariance matrices". In: *Journal of multivariate analysis* 88.2, pp. 365–411.

Miles, David and Jochen Schanz (2014). "The relevance or otherwise of the central bank's balance sheet". In: *Journal of International Economics* 92, S103–S116.

Mitchell, Mark, Todd Pulvino, and Erik Stafford (2004). "Price pressure around mergers". In: *The Journal of Finance* 59.1, pp. 31–63.

Modigliani, Franco and Richard Sutch (1966). "Innovations in interest rate policy". In: *The American Economic Review* 56.1/2, pp. 178–197.

Neely, Christopher J et al. (2010). *The large scale asset purchases had large international effects*. Federal Reserve Bank of St. Louis, Research Division.

Petajisto, Antti (2009). "Why do demand curves for stocks slope down?" In: *Journal of Financial and Quantitative Analysis* 44.5, pp. 1013–1044.

— (2011). "The index premium and its hidden cost for index funds". In: *Journal of Empirical Finance* 18.2, pp. 271–288.

Schnitzler, Jan (2016). "S&P 500 Inclusions and Stock Supply". In: *Working Paper*.

Scholes, Myron S (1972). "The market for securities: Substitution versus price pressure and the effects of information on share prices". In: *The Journal of Business* 45.2, pp. 179–211.

Shleifer, Andrei (1986a). "Do demand curves for stocks slope down?" In: *The Journal of Finance* 41.3, pp. 579–590.

— (1986b). "Do demand curves for stocks slope down?" In: *The Journal of Finance* 41.3, pp. 579–590.

Swanson, Eric T (2011). "Let's twist again: a high-frequency event-study analysis of operation twist and its implications for QE2". In: *Brookings Papers on Economic Activity* 2011.1, pp. 151–188.

Tobin, James (1969). "A general equilibrium approach to monetary theory". In: *Journal of money, credit and banking* 1.1, pp. 15–29.

Vayanos, Dimitri and Jean-Luc Vila (2009). *A preferred-habitat model of the term structure of interest rates*. National Bureau of Economic Research (No. w15487).

Wurgler, Jeffrey and Ekaterina Zhuravskaya (2002). "Does arbitrage flatten demand curves for stocks?" In: *The Journal of Business* 75.4, pp. 583–608.

Appendix

A Additional Material: Figures and Tables

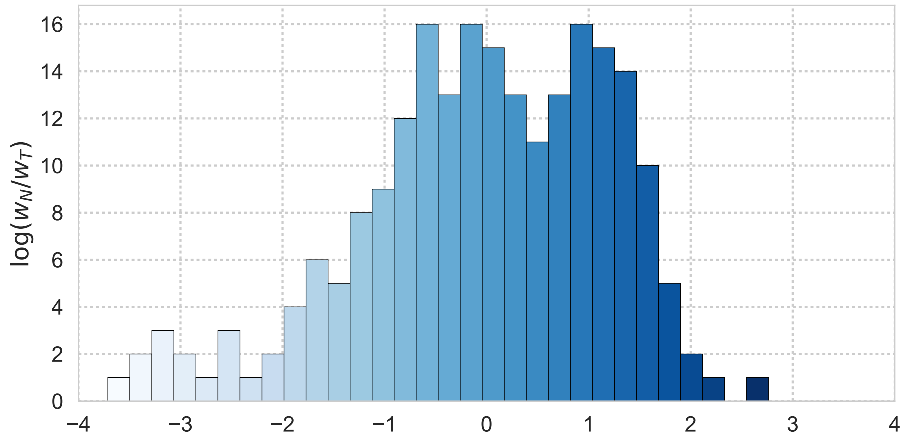

Figure A.1: Distortion. The figure plots the distribution of the log ratio between the Nikkei weight w_N and the TOPIX weight w_T for Nikkei firms only. The histogram shows a significant dispersion, confirming that Nikkei weights induce significant cross-sectional variation of purchased quantities relative to market capitalization.

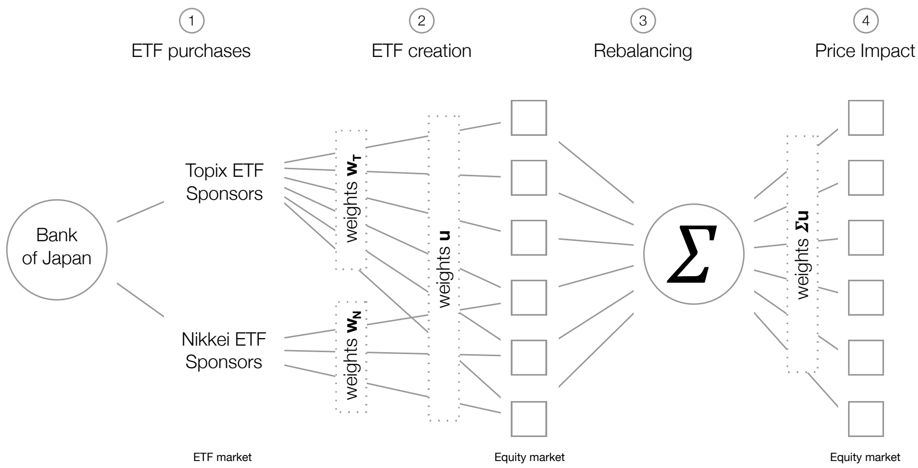

Figure A.2: From ETFs to equity This figure describes the channel though which ETF purchases of the central bank may have an impact on equity prices. As the BoJ buys TOPIX- and Nikkei-linked ETFs, these are created by ETF sponsors and/or authorized participants. The securities needed to form the ETF basket are collected by these intermediaries in the equity market, thus effectively reducing the supply of equity shares available to private investors.

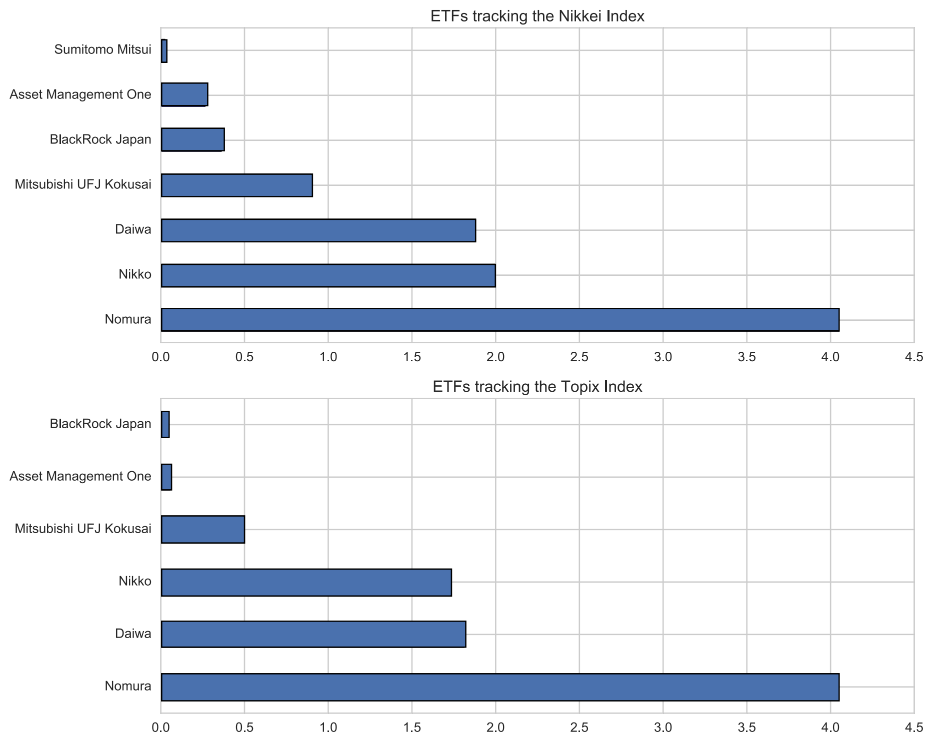

Figure A.3: Assets Under Management (AUM) by Provider (in trillion yen). This figure shows the Assets Under Management of ETFs aggregated at Provider level. The values are computed as of December 30th, 2016.

Panel A – 2014 Event

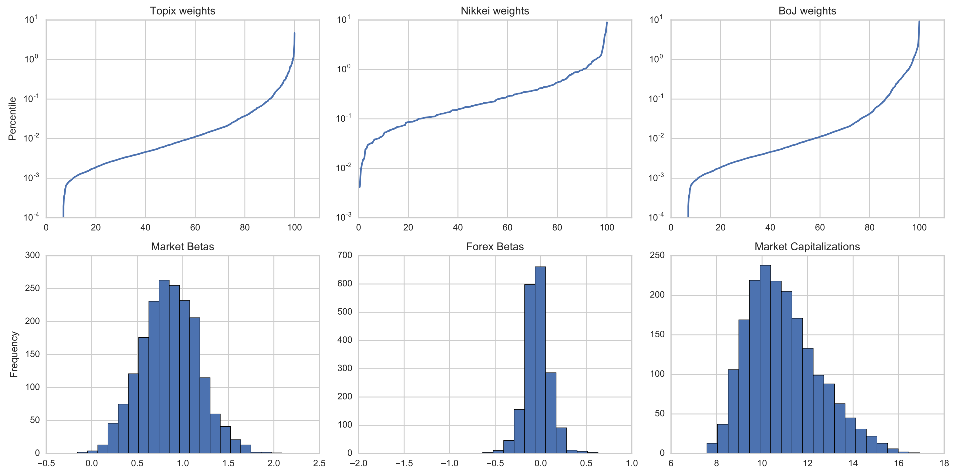

Panel B – 2016 Event

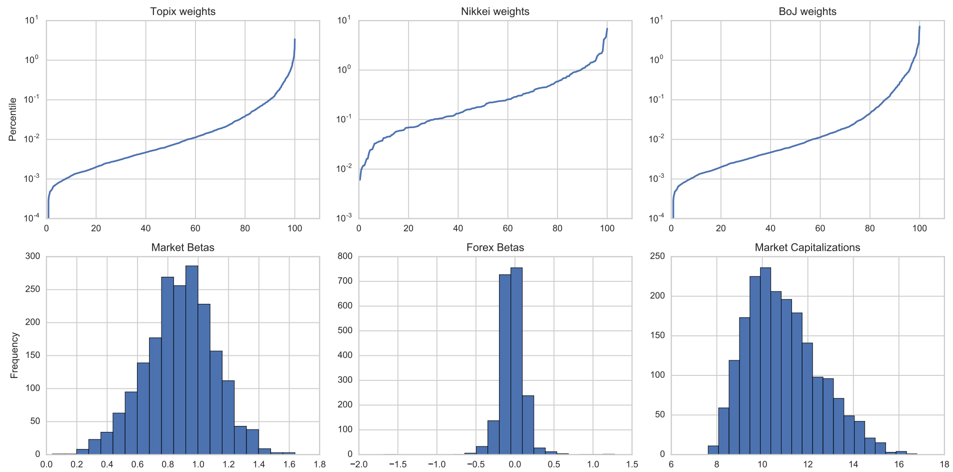

Figure A.4: Weights, betas and market capitalizations. The plots display cross-sectional heterogeneity of the variables of interest at the time of the BoJ announcements. Panel A refers to the announcement in 2014, Panel B to the announcement in 2016. The first row of each panel plots the percentile functions in logarithmic scale of the TOPIX weights (ω_T), the Nikkei weights (ω_N) and the BoJ weights(q). BoJ weights are computed as $\omega_T + \omega_N$ and correspond to the elements of the vector q in the model. The second row of each panel shows the distribution of stock-level market betas, Forex betas and market values. Market betas and Forex betas are estimated following the procedure explained in Section 3.5.3. Companies market capitalizations are in logs.

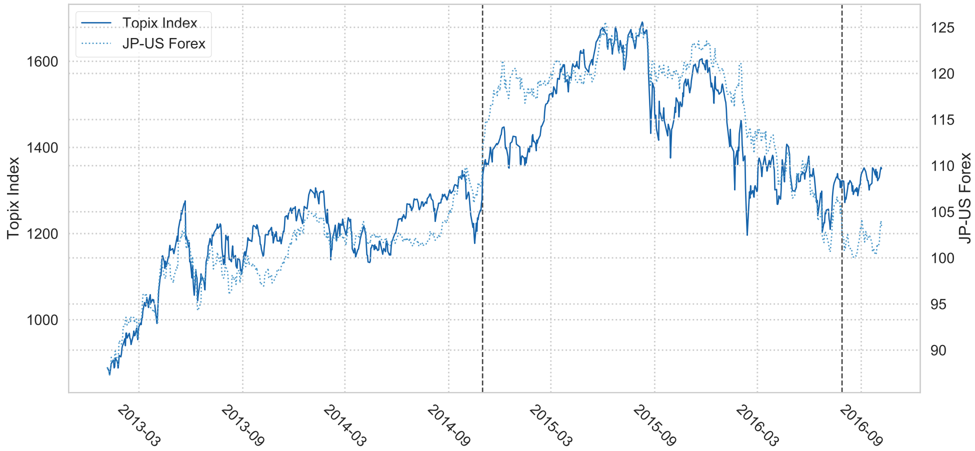

Figure A.5: TOPIX Index and JP-US Exchange Rate. This figure shows the time-series of the TOPIX Index over our sample period (green solid line, left axis) and of the exchange rate from US Dollar to Japanese Yen (purple dotted line, right axis).

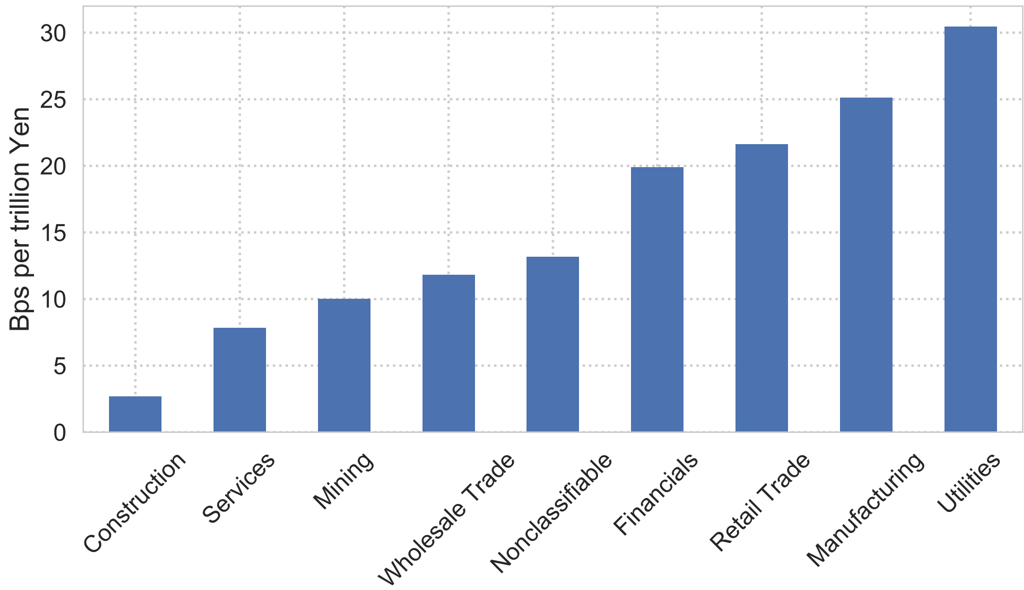

Figure A.6: Portfolio Balance Effect across Industries This figure shows the estimated portfolio balance impact of the policy, expressed in basis points per trillion Yen, computed separately for each sector.

	Mean	Std Deviation	Min	25%	50%	75%	Max	Obs
Market Cap (Billions Yen)								
TOPIX	252	853	2	17	45	148	22210	3824
Nikkei 225	1352	2110	28	294	683	1522	22210	442
Not Nikkei 225	108	251	2	15	35	94	4434	3382
Forex Beta								
TOPIX	-0.04	0.15	-1.68	-0.12	-0.04	0.04	1.27	3824
Nikkei 225	0.02	0.13	-0.39	-0.07	0.02	0.10	0.46	442
Not Nikkei 225	-0.05	0.15	-1.68	-0.12	-0.05	0.03	1.27	3382
Market Beta								
TOPIX	0.87	0.27	-0.16	0.69	0.88	1.05	2.08	3824
Nikkei 225	1.05	0.19	0.46	0.90	1.04	1.18	1.71	442
Not Nikkei 225	0.85	0.27	-0.16	0.67	0.85	1.02	2.08	3382
BoJ Weight								
TOPIX	0.05	0.22	0.00	0.00	0.00	0.01	4.65	3824
Nikkei 225	0.37	0.53	0.01	0.10	0.20	0.43	4.65	442
Not Nikkei 225	0.01	0.03	0.00	0.00	0.00	0.01	0.42	3382
Nikkei 225 Weight								
TOPIX	0.05	0.31	0.00	0.00	0.00	0.00	8.91	3824
Nikkei 225	0.45	0.82	0.00	0.09	0.20	0.45	8.91	442
Not Nikkei 225	0.00	0.00	0.00	0.00	0.00	0.00	0.00	3382
TOPIX Weight								
TOPIX	0.05	0.18	0.00	0.00	0.01	0.03	4.70	3824
Nikkei 225	0.30	0.44	0.01	0.06	0.15	0.36	4.70	442
Not Nikkei 225	0.02	0.05	0.00	0.00	0.01	0.02	0.85	3382

Table A.1: Summary Statistics. This table provides summary statistics for various stock characteristics by index membership. TOPIX stocks represent our entire sample of stocks. Nikkei stocks are those included in the Nikkei 225 index, while Not Nikkei stocks are those that only appear in the TOPIX index. All Nikkei companies also belong to the TOPIX index. All statistics are computed using pre-event information and pooling both events together.

Panel A: October 31st, 2014

	Raw Returns				Abnormal Returns			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
π	95.77***	91.16***	39.82***	38.19***	18.13***	13.14**	47.69***	43.85***
	(16.84)	(15.62)	(4.49)	(4.32)	(3.17)	(2.24)	(5.50)	(5.09)
u		0.01***	-0.00	-0.00		0.01***	-0.00	-0.00
		(3.32)	(-0.86)	(-0.73)		(3.58)	(-0.99)	(-0.77)
Market Beta			0.02***	0.01			-0.08***	-0.09***
			(2.70)	(1.38)			(-11.04)	(-12.29)
Forex Beta			0.07***	0.07***			0.07***	0.07***
			(7.58)	(7.45)			(8.30)	(7.88)
log(Market Cap)			0.01***	0.01***			0.01***	0.01***
			(6.87)	(6.33)			(7.13)	(6.60)
Amihud			0.00	0.00			0.00	0.00
			(0.96)	(0.51)			(0.94)	(0.50)
Observations	1,851	1,851	1,807	1,701	1,851	1,851	1,807	1,701
R-squared	0.13	0.14	0.19	0.24	0.01	0.01	0.11	0.17
Industry FE	NO	NO	NO	YES	NO	NO	NO	YES

Panel B: July 29th, 2016

	Raw Returns				Abnormal Returns			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
π	18.00***	16.27***	15.80***	18.59***	18.00***	16.64***	18.59***	20.24***
	(6.29)	(5.57)	(3.59)	(4.10)	(6.11)	(5.52)	(4.10)	(4.33)
u		0.01***	0.01*	0.01**		0.01**	0.01*	0.01**
		(2.75)	(1.82)	(2.00)		(2.10)	(1.80)	(1.97)
Market Beta			-0.00	-0.02*			-0.01	-0.02**
			(-0.57)	(-1.80)			(-0.99)	(-2.22)
Forex Beta			0.02**	0.01			0.02***	0.02*
			(2.06)	(1.42)			(2.63)	(1.94)
log(Market Cap)			0.00	0.00			-0.00	0.00
			(0.99)	(1.11)			(-0.02)	(0.17)
Amihud			0.00	0.00			0.00*	0.00*
			(1.48)	(1.53)			(1.75)	(1.78)
Observations	1,905	1,905	1,839	1,734	1,905	1,905	1,839	1,734
R-squared	0.02	0.02	0.03	0.05	0.02	0.02	0.03	0.05
Industry FE	NO	NO	NO	YES	NO	NO	NO	YES

Table A.2: Robustness: Alternative Covariance Matrix Estimation The tables report results for specifications similar to those in Table 3.2, but where the main explanatory variable $\pi = \Sigma u$ is constructed using the covariance matrix Σ estimated on raw returns. Regressions of event returns on the predicted price impact π are run separately for the two events. The dependent variable in columns 1-3 is the cumulative raw return, while in columns 4-6 is the cumulative abnormal return with respect to the market model estimated in the pre-event window. Cumulative returns are computed over a 10 days horizon after the announcement date. t-statistics are in parenthesis; asterisks denote conventional significance levels (***=1%, **=5%, *=10%).

	Abnormal Returns 2014						Abnormal Returns 2016					
Horizon (days)	5	10	21	63	126	252	5	10	21	63	126	252
π	11.06**	30.60***	22.51**	63.85***	152.1***	275.0***	15.52**	16.69**	30.48***	34.18**	82.20***	110.4***
	(2.11)	(4.62)	(2.52)	(4.34)	(8.50)	(10.32)	(2.96)	(2.52)	(3.42)	(2.32)	(4.60)	(4.14)
u	-0.00	-0.02***	-0.01	-0.02***	-0.02	-0.04*	0.010*	0.004	0.024**	0.014	0.045**	0.102***
	(-0.39)	(-3.20)	(-1.30)	(-2.52)	(-1.65)	(-1.99)	(2.06)	(0.67)	(2.55)	(1.35)	(2.91)	(4.70)
Market Beta	-0.04*	-0.07*	-0.08*	-0.17**	-0.32**	-0.45***	0.016	-0.01	0.009	0.012	0.007	0.041
	(-1.47)	(-1.97)	(-1.70)	(-2.32)	(-2.22)	(-2.68)	(0.59)	(-0.26)	(0.19)	(0.16)	(0.05)	(0.24)
Forex Beta	0.039**	0.041**	0.095***	0.087**	0.019	-0.12	-0.00	0.014	0.053**	0.049	0.190***	0.166
	(2.81)	(2.32)	(4.11)	(2.26)	(0.30)	(-0.03)	(-0.48)	(0.80)	(2.32)	(1.28)	(2.92)	(0.04)
log(Market Cap)	0.002	0.006	0.001	-0.00	0.005	-0.00	-0.00	-0.00	-0.01	-0.02	-0.06***	-0.10***
	(0.63)	(1.16)	(0.22)	(-0.07)	(0.26)	(-0.72)	(-1.40)	(-0.49)	(-1.60)	(-1.70)	(-3.43)	(-8.54)
Amihud	0.000	4.618	0.001	0.000	0.019***	0.005	0.000	0.000	-0.00	-0.00	-0.00**	-0.01***
	(0.66)	(0.03)	(0.75)	(0.36)	(7.19)	(1.67)	(0.18)	(0.11)	(-0.59)	(-1.31)	(-3.23)	(-3.11)
Observations	1,701	1,701	1,701	1,701	1,701	1,701	1,734	1,734	1,734	1,734	1,734	1,734
R-squared	0.114	0.160	0.102	0.120	0.180	0.141	0.071	0.050	0.101	0.111	0.203	0.191
Industry FE	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES

Table A.3: Cross-sectional regressions with industry fixed effects. The table report the coefficients of cross-sectional regressions of cumulative returns (in percentage points) computed at different horizons on the predicted price impact π and a set of control variables (standardized). In this specification we include industry fixed effects, based on the first 3 digits of the Standard Industry Classification Code (SIC-3). Regressions are run separately for the two events. The dependent variable is the cumulative abnormal return with respect to the market model estimated in the pre-event window. t-statistics from placebo regressions are in parenthesis; asterisks denote conventional significance levels (***=1%, **=5%, *=10%) based on empirical p-values.

	Abnormal Returns 2014						Abnormal Returns 2016					
	5	10	21	63	126	252	5	10	21	63	126	252
π	19.04***	42.17***	28.77***	72.38***	172.5***	300.7***	17.61**	17.34*	31.88***	38.51**	87.69***	112.9***
	(3.22)	(5.69)	(2.86)	(4.57)	(9.25)	(11.60)	(2.98)	(2.34)	(3.16)	(2.43)	(4.70)	(4.36)
u	0.008*	-0.00	-0.00	-0.03***	-0.03**	-0.11***	0.014***	0.001	-0.00	-0.02*	-0.04**	0.001
	(1.90)	(-1.32)	(-0.97)	(-3.44)	(-2.06)	(-4.64)	(3.28)	(0.27)	(-0.24)	(-2.15)	(-2.47)	(0.07)
Market Beta	-0.02	-0.05	-0.07	-0.16*	-0.29**	-0.40***	0.026	0.002	0.027	0.028	0.048	0.076
	(-0.84)	(-1.48)	(-1.37)	(-2.06)	(-2.05)	(-2.14)	(0.86)	(0.07)	(0.51)	(0.36)	(0.33)	(0.41)
Forex Beta	0.037**	0.042*	0.102***	0.097**	0.043	-0.13	-0.00	0.018	0.053*	0.060	0.196***	0.186***
	(2.20)	(2.02)	(3.81)	(2.28)	(0.61)	(-1.72)	(-0.18)	(0.90)	(2.00)	(1.43)	(2.73)	(2.33)
log(Market Cap)	0.001	0.006	0.001	0.001	0.005	-0.00	-0.00	-0.00	-0.01*	-0.03	-0.07***	-0.11***
	(0.33)	(0.94)	(0.16)	(0.06)	(0.24)	(-0.33)	(-1.48)	(-0.55)	(-1.80)	(-1.79)	(-3.58)	(-7.59)
Amihud	0.001	0.000	4.552	0.001	0.016***	0.018***	0.000	0.000	-0.00	-0.00	-0.00**	-0.00
	(0.79)	(0.25)	(0.00)	(0.67)	(6.13)	(7.14)	(0.48)	(0.39)	(-0.53)	(-0.96)	(-3.27)	(-2.45)
Nikkei	-0.01	-0.02*	-0.00	-0.00	-0.00	0.048*	-0.00	0.003	0.038*	0.049*	0.122***	0.153***
	(-1.26)	(-1.72)	(-0.27)	(-0.08)	(-0.18)	(1.54)	(-0.38)	(0.24)	(2.23)	(2.32)	(4.03)	(4.88)
Observations	1,807	1,807	1,807	1,807	1,807	1,807	1,839	1,839	1,839	1,839	1,839	1,839
R-squared	0.06	0.11	0.07	0.10	0.15	0.12	0.05	0.03	0.08	0.08	0.19	0.15
Industry FE	NO	NO	NO	NO	NO	NO	NO	NO	NO	NO	NO	NO

Table A.4: Cross-sectional regressions controlling for Nikkei. The table report the coefficients of cross-sectional regressions of cumulative returns (in percentage points) computed at different horizons on the predicted price impact π and a set of control variables (standardized). In this specification we add a dummy variable *Nikkei* that indicates stocks belonging to the Nikkei 225 Index. Regressions are run separately for the two events. The dependent variable is the cumulative abnormal return with respect to the market model estimated in the pre-event window. t-statistics from placebo regressions are in parenthesis; asterisks denote conventional significance levels (***=1%, **=5%, *=10%) based on empirical p-values.

	Abnormal Returns 2014						Abnormal Returns 2016					
Horizon (days)	5	10	21	63	126	252	5	10	21	63	126	252
π	12.21**	32.64***	22.95**	63.97***	150.7***	266.7***	15.62**	16.37**	28.39**	31.48**	75.75***	102.7***
	(2.28)	(4.80)	(2.51)	(4.38)	(8.72)	(10.01)	(2.91)	(2.41)	(3.11)	(2.16)	(4.38)	(3.86)
u	0.006	-0.00	-0.00	-0.02***	-0.03**	-0.10***	0.012**	0.000	-0.00	-0.02*	-0.04**	-0.00
	(1.35)	(-1.41)	(-1.16)	(-2.89)	(-2.32)	(-4.28)	(2.72)	(0.05)	(-0.56)	(-2.42)	(-2.95)	(-0.21)
Market Beta	-0.04*	-0.07*	-0.08*	-0.17**	-0.33**	-0.46***	0.016	-0.00	0.011	0.015	0.014	0.050
	(-1.48)	(-1.98)	(-1.71)	(-2.38)	(-2.23)	(-2.28)	(0.59)	(-0.25)	(0.24)	(0.21)	(0.10)	(0.25)
Forex Beta	0.040**	0.043**	0.095***	0.087**	0.018	-0.12***	-0.00	0.012	0.044*	0.038	0.163**	0.134***
	(2.83)	(2.38)	(4.11)	(2.25)	(0.28)	(-1.93)	(-0.45)	(0.71)	(1.93)	(0.98)	(2.49)	(2.01)
log(Market Cap)	0.002	0.006	0.001	-0.00	0.004	-0.00	-0.00	-0.00	-0.01	-0.02	-0.06***	-0.10***
	(0.61)	(1.09)	(0.21)	(-0.07)	(0.24)	(-0.81)	(-1.29)	(-0.48)	(-1.61)	(-1.69)	(-3.43)	(-9.10)
Amihud	0.001	0.000	0.001	0.000	0.019***	0.005	0.000	0.000	-0.00	-0.00	-0.00**	-0.01***
	(0.67)	(0.06)	(0.70)	(0.34)	(6.97)	(1.52)	(0.18)	(0.09)	(-0.64)	(-1.32)	(-3.33)	(-3.15)
Nikkei	-0.01	-0.01	-0.00	-0.00	0.011	0.074***	-0.00	0.006	0.039*	0.050*	0.120***	0.143***
	(-1.04)	(-1.39)	(-0.22)	(-0.05)	(0.36)	(2.15)	(-0.17)	(0.46)	(2.24)	(2.29)	(3.62)	(4.17)
Observations	1,701	1,701	1,701	1,701	1,701	1,701	1,734	1,734	1,734	1,734	1,734	1,734
R-squared	0.12	0.16	0.10	0.12	0.18	0.14	0.07	0.05	0.11	0.12	0.21	0.20
Industry FE	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES

Table A.5: Cross-sectional regressions controlling for Nikkei and industry. The table report the coefficients of cross-sectional regressions of cumulative returns (in percentage points) computed at different horizons on the predicted price impact π and a set of control variables (standardized). In this specification we add a dummy variable *Nikkei* that indicates stocks belonging to the Nikkei 225 Index and industry fixed effects based on the first 3 digits of the Standard Industry Classification Code (SIC-3). Regressions are run separately for the two events. The dependent variable is the cumulative abnormal return with respect to the market model estimated in the pre-event window. t-statistics from placebo regressions are in parenthesis; asterisks denote conventional significance levels (***=1%, **=5%, *=10%) based on empirical p-values.

B Model Derivation

The model features a representative investor who chooses time-t demand N_t of shares to maximize its next period exponential utility subject to a standard budget constraint

$$\max_{N} \; \mathrm{E}_t\left(-\exp(-\gamma W_{t+1})\right) \tag{32}$$

$$\text{s.t.} \qquad W_{t+1} = W_t(1+r) + N_t'(p_{t+1} + D_{t+1} - p_t(1+r)) \tag{33}$$

From the first order condition it follows that

$$N_t = \frac{1}{\gamma}[\mathrm{Var}_t(p_{t+1} + D_{t+1})]^{-1}(\mathrm{E}_t[p_{t+1} + D_{t+1} - p_t(1+r)]) \tag{34}$$

We restrict our attention to the covariance stationary equilibrium. Imposing market clearing and substituting $V = \mathrm{Var}_t(p_{t+1} + D_{t+1})$ yields

$$(1+r)p_t = \mathrm{E}_t[p_{t+1} + D_{t+1}] - \gamma V Q_t \tag{35}$$

Iterating forward up to time T and applying the law of iterated expectations we get

$$(1+r)p_t = \mathrm{E}_t\left[\frac{p_T}{(1+r)^{T-t-1}}\right] + \sum_{i=0}^{T-t-1} \frac{D_t}{(1+r)^i} - \gamma V \sum_{i=0}^{T-t-1} \frac{\mathrm{E}_t[Q_{t+i}]}{(1+r)^i} \tag{36}$$

Taking the limit $T \to \infty$ and imposing the no-bubble condition yields

$$p_t = \frac{D_t}{r} - \frac{\gamma V}{(1+r)}\left(\sum_{i=0}^{\infty} \frac{\mathrm{E}_t[Q_{t+i}]}{(1+r)^i}\right) = \frac{1}{r}\left(D_t - \gamma V \Omega_t\right) \tag{37}$$

where we introduced the notation

$$\Omega_t = \frac{r}{1+r} \sum_{i=0}^{\infty} \frac{\mathrm{E}_t[Q_{t+i}]}{(1+r)^i} \tag{38}$$

which can be interpreted as the discounted time-t expected future supply of the assets. This term is crucial for our analysis, representing the channel through which the central bank is able to affect risk premia. Under no expectation of monetary policy intervention we have $\mathrm{E}_t(Q_{t+i}) = Q$, so that the resulting pricing equation collapses to

$$p_t = \frac{1}{r}\left(D_t - \gamma V Q\right) \tag{39}$$

where the vector $\gamma V Q$ can be interpreted as the cross-sectional vector of risk premia required by investors in equilibrium. In our context, this is the pricing equation that applies before the policy announcement at $t = 1$.

In the following sections we look at what happens to prices if the central bank unexpectedly commits itself to a large-scale purchase of assets over a defined period, thus affecting the expected path of future supply Ω_t.

We now solve the model in its most general form, allowing for the possibility that agents expectations on future supply change over time. We assume that for each $t \geq 1$ there exist a scalar $\lambda_t \geq 0$ such that the time-t expectation is

$$\begin{cases} \mathrm{E}_t(Q_{t+i}) = Q & \text{for } i \geq 0 \text{ and } t < 1 \\ \mathrm{E}_t(Q_{t+i}) = Q - \lambda_t(t+i)q & \text{for } i \geq 0 \text{ and } t = 1, \dots, M \\ \mathrm{E}_t(Q_{t+i}) = Q - \lambda_t M q & \text{for } i \geq M - t \text{ and } t \geq 1 \end{cases} \tag{40}$$

The parameter λ_t can be interpreted as the degree of confidence of investors in the BoJ commitment or, in other words, as the conditional probability they attach to the continuation of the program.

Assuming that investors increase their confidence as time passes – and they observe more actual purchases by the BoJ – amounts to assume that λ_t is increasing in time.

After the BoJ announcement, for $t \geq 1$, the expected supply can be written as

$$\begin{aligned}
\Omega_t &= \frac{r}{1+r}\sum_{i=0}^{\infty}\frac{\mathrm{E}_t[Q_{t+i}]}{(1+r)^i} && (41)\\
&= \frac{r}{1+r}\left(\sum_{i=0}^{M-t-1}\frac{Q-\lambda_t(t+i)q}{(1+r)^i} + \sum_{i=M-t}^{\infty}\frac{Q-\lambda_t Mq}{(1+r)^i}\right) && (42)\\
&= Q - \frac{\lambda_t r}{1+r}\left(\sum_{i=0}^{M-t-1}\frac{(t+i)q}{(1+r)^i} + \sum_{i=M-t}^{\infty}\frac{Mq}{(1+r)^i}\right) && (43)\\
&= Q - \frac{\lambda_t r}{1+r}\left(\sum_{i=0}^{M-t-1}\frac{(t+i-M)q}{(1+r)^i} + \sum_{i=0}^{\infty}\frac{Mq}{(1+r)^i}\right) && (44)\\
&= Q - \lambda_t Mq + \frac{\lambda_t r}{1+r}\sum_{i=0}^{M-t-1}\frac{(M-t-i)}{(1+r)^i}q && (45)\\
&= Q - \lambda_t Mq + \lambda_t \varphi(t) q && (46)
\end{aligned}$$

where we introduced the real-valued function

$$\varphi(t) = \frac{r}{1+r}\sum_{i=0}^{M-t-1}\frac{(M-t-i)}{(1+r)^i}, \qquad t \geq 1 \tag{47}$$

This quantity represents the residual duration of the program at time t. In the Internet Appendix we show that, for realistic values of the risk-free rate r, the function $\varphi(t)$ enjoys the following properties:

(i) $\varphi(t+1) - \varphi(t) < 0$

(ii) $\varphi(t) < M$ for $t \geq 1$

(iii) $\varphi(t) = 0$ for $t \geq M$

Given the assumption that by the end of the policy horizon $t = M$ the central bank will have purchased exactly Mq as announced and afterwards it will not engage in further market operations, it follows that the pricing equation (39) takes the form

$$\begin{cases} p_t = \frac{1}{r}\left(D_t - \gamma V Q\right) & \text{for } t < 1 \\ p_t = \frac{1}{r}\left(D_t - \gamma V (Q - \lambda_t M q + \lambda_t \varphi(t) q)\right) & \text{for } t = 1, \dots M \\ p_t = \frac{1}{r}\left(D_t - \gamma V (Q - M q)\right) & \text{for } t \geq M \end{cases} \tag{48}$$

and the price change at the announcement day $t = 1$ can be written as

$$p_1 - p_0 = \frac{1}{r}\left(\varepsilon_1 + \lambda_1 \gamma V (Mq - \varphi(1) q)\right) \tag{49}$$

Dividing by p_0 coordinate-wise proves Proposition 1. The equation also shows that the size of the price jump is increasing in the initial belief parameter λ_1.

On the days following the announcement, price changes depend on the time-series evolution of λ_t. Denoting the updates in beliefs by $\Delta\lambda_{t+1} = \lambda_{t+1} - \lambda_t$ we have

$$\begin{aligned} p_{t+1} - p_t &= \frac{1}{r}\left(\varepsilon_{t+1} - \gamma V ((\lambda_{t+1}\varphi(t+1) - \lambda_t \varphi(t)) - \Delta\lambda_{t+1} M) q\right) && (50) \\ &= \frac{1}{r}\left(\varepsilon_{t+1} + \gamma \xi(t+1) V q\right), \qquad t = 1, \dots M && (51) \end{aligned}$$

Given $\Delta\lambda_{t+1} > 0$, the following inequalities show that $\xi(t) > 0$

$$\lambda_{t+1}\varphi(t+1) - \lambda_t\varphi(t) < \lambda_{t+1}\varphi(t) - \lambda_t\varphi(t)) = \Delta\lambda_{t+1}\varphi(t) < \Delta\lambda_{t+1} M \tag{52}$$

Therefore we conclude that if $\Delta\lambda_{t+1} > 0$ for every $t = 1, \dots, M$, then we should observe a positive relationship between Vq and the cross-section of price changes. To complete the proof of the first part of Proposition 2 we need to show that this conclusion also applies to the relationship between returns $R_{i,t+1} = (p_{i,t+1} - p_{i,t})/p_{i,t}$ and $\pi = \Sigma u$. This follows from the definitions of $\Sigma_{i,j}$, u_i and π_i

$$R_{i,t+1} = \frac{1}{r}\left(\varepsilon_{i,t+1}/p_{i,t} + \gamma\xi(t+1)(Vq)_i/p_{i,t}\right) = \frac{1}{r}\left(\varepsilon_{i,t+1}/p_{i,t} + \gamma\xi(t+1)\pi_i\right) \tag{53}$$

Finally taking the expectation of the cumulative returns we get

$$\sum_{s=1}^{t} \mathrm{E}[R_s] = \sum_{s=1}^{t} \frac{1}{r}\left(\gamma\xi(s)\pi\right) = \theta_t \pi \tag{54}$$

where $\theta_t = \sum_{s=1}^{t} \frac{\gamma}{r}\xi(s)$ is a positive and increasing function of t, which follows from $\xi(s) > 0$ for $s = 1, \dots, M$ as shown above. This concludes the proof of Proposition 2.

C Systematic Risk in the Model

In our model the systematic risk of security i is measured as $(VQ)_i$, where V is the covariance matrix of price innovations and Q is the vector of shares outstanding. This quantity represents the covariance of the security's price changes ε with the change in the wealth of the representative agent (i.e. the value of the market portfolio) and it therefore admits an interpretation similar to the market beta. Denoting the value of the market portfolio by MP and the covariance of the price of stock i with MP by $\beta(P)_i$ we have

$$\beta(P)_i = \mathrm{Cov}(\Delta MP, \Delta p_i) \propto \mathrm{Cov}(\varepsilon_t' Q, \varepsilon_{i,t}) = (VQ)_i \tag{55}$$

Market betas are usually defined in terms of returns, not of price changes. Thus an empirically more relevant definition of the systematic risk of security i is given by $(\Sigma W)_i$, where Σ is the covariance matrix of returns and W is the vector of percentage weights of the market portfolio. This quantity is proportional to the market beta of stock i, denoted by $\beta(R)_i$

$$\beta(R)_i = \frac{\mathrm{Cov}(R^{mkt}, R_i)}{\mathrm{Var}(R^{mkt})} \propto \mathrm{Cov}(R'W, R_i) = \sum_{i,j} \mathrm{Cov}(R_j, R_i) W_j = (\Sigma W)_i \tag{56}$$

Let $Mq = Q^{\text{post}} - Q$ denote the announced change in the supply of assets and β^{post} the implied vector of market beta after the announcement. It follows immediately from the above definitions that the change in (price-level) systematic risk is proportional to the product between V and q:

$$\beta(P)_i^{\text{post}} - \beta(P)_i \propto (VQ^{\text{post}} - VQ)_i \propto -(Vq)_i \tag{57}$$

Similarly, from the definition of $\pi = \Sigma u$, where $u_i = p_i q_i$ is the announced change in the supply of stock i expressed in yen, it follows that

$$\beta(R)_i^{\text{post}} - \beta(R)_i \propto (\Sigma W^{\text{post}} - \Sigma W)_i \propto -(\Sigma u)_i = -\pi_i \tag{58}$$

For each stock i, π_i can thus be interpreted as the change in the stock beta (i.e. the systematic risk) induced by the supply shock. Notice that this change is induced by the policy through a modification of the portfolio held by the representative agent, while the fundamental covariance structure of returns is unchanged.