FINNISH CENTRE OF EXCELLENCE IN TAX SYSTEMS RESEARCH

FIT Working Paper 45

Risto Hurmeranta and Teemu Lyytikäinen

Nominal Loss Aversion in the Housing Market and Household Mobility

Tampere University
UNIVERSITY OF HELSINKI
VATT INSTITUTE FOR ECONOMIC RESEARCH
CENTRES OF EXCELLENCE IN RESEARCH

Nominal Loss Aversion in the Housing Market and Household Mobility

Risto Hurmeranta and Teemu Lyytikäinen *

Abstract

Households are averse to realizing nominal housing market losses. Reduced household mobility, in an attempt to avoid selling at a loss, implies misallocation of housing and can affect the functioning of the labor market. However, direct evidence on mobility responses is scarce. This paper studies the effect of expected losses on homeowners' propensity to move using administrative data on housing transactions matched with detailed data on household characteristics. We use an ensemble machine learning method to estimate expected prices for the universe of apartments in the three largest travel-to-work areas in Finland in 2006–2018. We find that homeowners below the zero-return cutoff are roughly 44% less likely to move than those above the cutoff. The effect of loss aversion on sales is larger than the effect on mobility. Loss-exposed homeowners are more likely to move without selling their previous home. Conversion of previous apartment to rental seems to enable homeowners to move without realizing nominal losses. They are also more likely to become renters themselves. Expected losses also reduce inter-regional mobility, which suggests that loss aversion can lead to misallocation of the labor force, but the impact on labor market transitions appears to be economically small.

JEL: R21; R23; R31

Keywords: Nominal loss aversion, housing market, household mobility.

*Hurmeranta: University of Helsinki and Finnish Centre of Excellence in Tax Systems Research (FIT) (e-mail: risto.hurmeranta@helsinki.fi); Lyytikäinen: VATT Institute for Economic Research and Finnish Centre of Excellence in Tax Systems Research (FIT) (e-mail: teemu.lyytikainen@vatt.fi). We thank Lu Han, Tuukka Saarimaa and Pablo Warnes and seminar participants at the European Meeting of the Urban Economics Association, Aalto University, FIT and VATT Institute for Economic Research for helpful comments and discussions. Lyytikäinen acknowledges funding from the Research Council of Finland (grant no. 345384, 370060 and 364189). Hurmeranta acknowledges funding from the Yrjö Jahnsson Foundation (grant no. 20237631 and 20237657), the Kone Foundation (grant no. 202308821), and Research Council of Finland (grant no. 359536).

1 Introduction

Household mobility is essential for the economy to be able to adjust to regional shocks and for efficient allocation of the housing stock to residents within regions. Mobility rates are substantially lower among homeowners than renters. Understanding the behavioral mechanisms behind this difference is highly important, as in most OECD countries more than two-thirds of households own their home.[1]

Nominal loss aversion among homeowners is a potential explanation of the low mobility rates of homeowners, along with transaction costs[2] and sorting. Nominal loss aversion was first documented in the housing market context by Genesove and Mayer (2001), who draw on the prospect theory of Kahneman and Tversky (1979), and argue that homeowners have a tendency to anchor their pricing decisions on the nominal price they had paid for the apartment. Losses relative to this reference point reduce utility more than equally large gains increase it. Consistent with loss aversion, they find that list prices conditional on characteristics are higher among sellers facing a loss.

Subsequent empirical literature confirms that expected losses lead to higher listing and sales prices conditional on housing characteristics (Anenberg, 2011; Ross and Zhou, 2021; Bracke and Tenreyro, 2021; Greenaway-McGrevy and Sorensen, 2021), and result in a longer time on the market (Bokhari and Geltner, 2011) and a lower propensity to sell (Einiö et al., 2008). Andersen et al. (2022) provide the first structural model that embeds loss aversion in households' sales decisions. Bokhari and Geltner (2011) demonstrate that incorporating the implications of loss aversion into housing market indices can also enhance their precision.

Loss aversion can also reduce the mobility of homeowners, leading to a misallocation of the housing stock, and possibly the labor force. However, direct evidence of the implications of loss aversion for the likelihood of moving is limited. The few exceptions include Engelhardt (2003) and Steegmans and Hassink (2018), and these are consistent with nominal loss aversion reducing mobility.

This paper studies the effect of homeowners' prospective losses on the propensity to move and the propensity to sell using administrative data on the stock of Finnish apartments and transactions matched with data on the owners and residents. Our data show that transaction volume is not a synonym for mobility: roughly 20% of homeowners who move do not sell their previous unit. We study whether loss aversion differentially affects selling and moving, and analyze interesting behavioral margins through which homeowners facing a loss may avoid selling when they move. Moreover, we study whether different types of moves are differentially affected, to gain insights into the broader implications

[1] OECD (2024), OECD Affordable Housing Database - indicator HM1.3 Housing tenures, https://oe.cd/ahd

[2] There is strong evidence that real estate transfer taxes reduce mobility of homeowners (See e.g. Eerola et al., 2021).

of loss aversion for housing markets and labor markets.

Our empirical strategy is to first estimate the predicted market value of each apartment in each year. Then we plot homeowners' mobility rates in bins of expected rate of nominal returns. Consistent with loss aversion, visual examination of the binned mobility rates indicates that an expected loss leads to a clear drop in mobility. The effect on mobility is quantified by estimating models of mobility on a dummy for expected nominal loss, while flexibly controlling for the expected rate of nominal returns. Furthermore, we control for household leverage[3] and the duration of the residence spell.

Accurate prediction of the market value is important for the feasibility of the empirical strategy. We use ensemble machine learning and find that it vastly outperforms the traditional hedonic model estimated with OLS. Close to the zero-gain cutoff, we may still erroneously classify apartments into the loss and gain domains, or owners may be uncertain as to whether to expect a loss or a gain. To address this issue, we drop observations very close to the cutoff when quantifying the effect.

The broader economic implications of loss aversion depend on whether reduced mobility is limited to adjustments in housing consumption within the city or extends to a longer-distance relocation of the labor force, but there is no prior evidence on how loss aversion affects different types of moves. Since Oswald (1996), many studies have analyzed whether homeownership can impede the mobility of the labor force and lead to higher unemployment through regional mismatch in the labor market. Evidence on the link between homeownership and labour market outcomes is mixed and inconsistent, partly because the cause and effect are difficult to disentangle (Goodman and Mayer, 2018). We provide new insights into this literature by examining differential impacts of loss aversion on labor market moves and housing-related moves. In addition, we analyze the propensity to change job or to become employed if unemployed.

We find that homeowners below the zero-return cutoff are roughly 50% less likely to sell than those above the cutoff. The effect of loss aversion on mobility is smaller, roughly 44%. This is possible because homeowners with an expected loss are more likely to move without selling their previous apartment. Renting out the previous apartment without selling seems to allow some homeowners to avoid realizing nominal losses when moving. Moreover, we find that homeowners facing a loss are more likely to become renters conditional on moving. These results imply that loss aversion can tilt the composition of households' investment portfolios and shape tenure type distribution in the housing market.

We find that expected losses reduce both intra-regional and inter-regional mobility, which implies that loss aversion leads to misallocation of the housing stock and can also lead to misallocation of the labor force across regions. Our analysis of labor market

[3]Expected losses may trigger down-payment constraints that reduce sales volume (Stein, 1995). Ferreira et al. (2010) find that negative equity affects mobility.

transitions suggests that reduced inter-regional mobility leads to fewer job changes, but the effect appears economically small because homeowners' inter-regional mobility is low in the baseline. There is no clear evidence on the reduced flow from unemployment.

These novel findings on how behavioral biases affect homeowners' decision-making are highly important, as housing markets have a significant impact on the economy as a whole. The behavior of owner-occupier households is of particular interest for policy design, as homeownership is heavily subsidized through the tax code in most countries.[4]

This paper is the first study on nominal loss aversion to separate selling from moving, to document the role of conversion of old apartments to rental as an escape-valve, and to analyze labor market implications. In addition, we contribute to the literature on how housing market conditions and policies affect mobility (e.g. Fonseca and Liu, 2024; Eerola et al., 2021; Ihlanfeldt, 2011; Ferreira et al., 2010). We also contribute to the broader behavioral economics literature (e.g. Thaler, 2016). Recent studies have debated the relevance of loss aversion (Oprea, 2024; Banki et al., 2025). Our findings provide additional evidence that loss aversion affects some of the most important economic decisions households make: whether to sell a house and whether to move.

The rest of the paper is organized as follows. Section 2 describes the institutional setting and lays out testable hypotheses. Section 3 introduces the data. Section 4 presents the empirical strategy and estimation methods. Section 5 presents the results and Section 6 concludes.

2 Institutional setting and hypotheses

2.1 Institutional setting

Homeownership is the dominant tenure type in Finland, with roughly two-thirds of households owning their home.[5] Approximately 45% of homeowners own an apartment and 55% own a single family home. We focus on apartment owners as our data lack transactions in directly owned single family homes. The ownership of apartments is organized through non-profit housing companies that can also be characterized as cooperatives of owners.[6] The owner of the apartment owns shares in the housing company that give the owner full control over the apartment. The owner can occupy the apartment herself or rent it out freely, which means that individual apartments can flexibly switch between

[4] According to estimates by the Finnish Ministry of Finance, the loss of tax revenue due to the non-taxation of the imputed rental income and capital gains of homeowners amounted to 5.8 bn euros or 2.3 percent of GDP in 2021.

[5] The remaining third of households are renters in the unregulated private rental sector or in the equally large social housing sector.

[6] Housing companies typically own multi-unit buildings like apartment blocks or row-houses, but sometimes also single family detached units are owned through a housing company. For brevity, we refer to all housing units owned through a housing company as apartments.

owner-occupied and rental sectors.

The buyer of the apartment can take out a personal housing loan using the apartment as collateral. Housing loans are transferable, and roughly 95% of buyers choose a variable interest rate, which means that falling interest rates do not cause a mortgage lock-in[7]. Loan duration is typically 20–30 years. Interest expenses were fully deductible in taxation until 2010, after which the deductible share gradually decreased to 35% in 2018, the last year of our analysis period. Deductibility was phased out completely in 2023. Housing companies often have outstanding debt, which is typically allocated to shareholders according to the floor area of the apartments. The owner can pay down her share of the housing company debt, or service the debt through a monthly payment to the housing company.

Apartment transactions are subject to a transfer tax. Before 2013, the transfer tax rate was 1.6% and it applied to the transaction price. In March 2013, the transfer tax rate was raised to 2% and the tax base was extended to include the outstanding housing company debt attached to the apartment. Capital gains on owner-occupied housing are tax-free after two years of continuous occupation of the unit. Capital gains on investment apartments are taxed at the capital income tax rate. Municipalities levy property taxes, but there are no caps on tax increases that would discourage mobility, unlike in some US states (see e.g. Ihlanfeldt, 2011).

The period we study is characterized by overall stagnant development in real housing prices in Finland. We focus on travel-to-work areas (TWA) around the three largest urban centers in Finland: Helsinki, Tampere, and Turku, where housing prices have grown both in nominal and real terms. We focus on these regions because they have a reasonably dense apartment market which improves the precision of our prediction model. Figure 1 shows the nominal housing price index for these regions constructed by the authors from the transaction data. The Helsinki TWA has a higher price level than Tampere and Turku. In all three TWAs, housing prices increased steadily in nominal terms. Despite these trends, in Section 3 we find that for approximately 10% of the apartment-year observations in our data, the predicted nominal market value was lower than the previous purchase price.

[7]When interest rates fall, owners with non-assumable fixed rate mortgages forgo moves to keep low mortgage rates (Fonseca and Liu, 2024; Quigley, 1987).

Figure 1: Nominal housing price index (Tampere TWA in 2006=1)

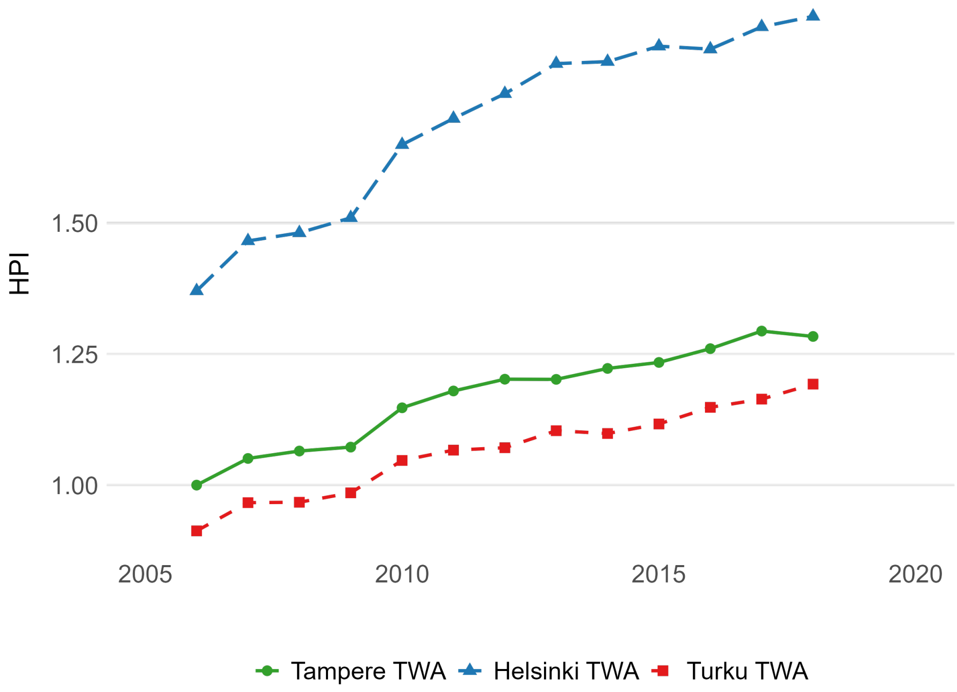

Notes: The figure plots the estimated nominal housing price index for each travel-to-work area (TWA) over the years 2006–2018. The index is based on a hedonic price regression controlling for apartment characteristics, age, and distance to the city center, with year and TWA fixed effects and their interactions. The index is normalized to the Tampere TWA in 2006.

2.2 Hypotheses

The previous literature starting from Genesove and Mayer (2001) has established that loss aversion and reference dependence affect sellers' behavior in the housing market. If homeowners value losses higher than equally large gains, there is a sharp increase in the optimal listing price premium when the expected sales price falls below the initial purchase price (e.g. Andersen et al., 2022). A sharp increase in the listing premium should lead to a sharp reduction in interested buyers and the propensity to sell. In addition, households expecting not to get offers matching their purchase price may refrain from listing altogether. Both mechanisms should lead to a reduction in homeowners' mobility. Thus, our first hypothesis is that there should be an abnormal drop in moving propensity close to the zero nominal gain cutoff so that homeowners below the cutoff are less likely to move.

In addition to a drop in mobility at zero, reference dependence and nominal loss aversion can affect the functional form of the relationship between expected gain and moving propensity in other parts of the distribution. Our empirical analysis focuses solely on the drop in mobility at the cutoff, which can be detected through transparent graphs plotting moving propensity over the distribution of expected gains, and credibly attributed to loss aversion, as there is no other reason to expect a discontinuous change at the zero gain cutoff.

We would expect a drop in both sales propensity and moving propensity around the zero gain cutoff, but these effects may differ, because it is also possible to move without selling. According to Table 1 in Section 3, moving without selling the previous unit seems to be quite common in our data. The table shows that roughly one-fifth of homeowners who move do not sell their previous apartment by the end of the following year.

In the Finnish setting, the fact that apartments can be easily rented out at unregulated market rents means that owner-occupiers considering moving have the option to rent out their previous apartment rather than sell it. We expect this to be a particularly attractive alternative for an owner who faces a nominal loss upon sale. This leads to our second hypothesis that expecting a nominal loss upon sale has a greater impact on the propensity to sell than on the propensity to move, and, conditional on moving, those facing a loss are more likely to rent out their previous unit.

Naturally, homeowners who move without selling may be credit-constrained when buying a new apartment, but a well-functioning rental market would mean that they will nevertheless be able to move by becoming renters. The third hypothesis we test is that mobility to rental housing is less strongly affected by loss aversion than is mobility to owner-occupied housing.

Our fourth hypothesis pertains to differential effects of loss aversion by distance of move. Moves within a TWA are probably often motivated by adjustments in housing consumption (e.g. upsizing or downsizing by one room), while inter-TWA moves include more momentous shifts in the economic situation of the household, like a change to a new job in another region or major life events (e.g. relocation after retirement). The disutility of selling at a loss may affect different kinds of moves differently. If the benefits of within-TWA moves are more often incremental compared to the current unit, we would expect within-TWA moves to be more strongly affected by loss aversion than between-TWA moves. Testing for differential effects of loss aversion on within- and between-TWA moves is of interest, because it provides insights into whether loss aversion can lead to misallocation of the labor force across regions, in addition to misallocation of housing within the region.[8]

[8]Real estate transfer taxes have been found to discourage short-distance and housing-related moves more strongly than longer-distance moves, which are more likely to be motivated by labor market considerations (Eerola et al., 2021; Hilber and Lyytikäinen, 2017). We test whether this also applies to loss aversion.

3 Data

Our data come from Statistics Finland and cover all market transactions and ownership spells for dwellings in apartment buildings and row houses in Finland between 2006-2018. The data include key characteristics of apartments, including the number of rooms, floor area, and location (250m x 250m grid cells). We also observe key characteristics of the owners such as disposable income, housing debt balance, and end-of-year residence.[9]

For the analysis, we compile two separate data sets. The matched apartment-household panel tracks homeowners' residence spells and is used in the main analysis on mobility. The transaction data set is used to build a hedonic pricing model for apartment market values that are used to calculate homeowners' expected nominal returns.

We focus on the mobility of owner-occupier households living in the three largest travel-to-work areas in Finland around the cities of Helsinki, Turku and Tampere.[10] The main outcomes on mobility and sales are based on changes in end-of-year residences and owners. If two owners move to separate addresses, we define these moves as undefined with respect to ownership and the location of the new apartment. In cases with two owners, we require both to move in order to define the household as mover.

Table 1 presents the summary statistics for the matched apartment-household panel. The price model predicts roughly 10% of observations expecting a nominal loss. The mean propensity to sell is 5.8% and the mobility rate is 6.9%. The higher mobility rate is possible because 17% of movers do not sell their previous apartment by the end of the following year. 12% of movers become landlords, and 12% of movers become renters themselves. Intra-regional mobility is much more common than inter-regional.

[9]See Appendix A1.1 for a full description of the data sources and sample restrictions. Replication code for data preprocessing and statistical analysis is available from the authors upon request.

[10]As defined by Statistics Finland, a TWA is formed by a central municipality and surrounding municipalities from which at least 10 percent of the labor force commute to the central municipality. Central municipalities are municipalities from which at most 25 percent commute to other municipalities.

Table 1: Summary statistics

	mean	std	p01	p99
Key explanatory variables				
Expected return	0.866	6.956	-0.155	6.075
Spell length	10.595	7.452	1.216	29.956
LTV	0.290	0.354	0.000	1.143
Loss	0.104	0.305	0.000	1.000
Gain	0.896	0.305	0.000	1.000
Outcome				
Sale	0.058	0.234	0.000	1.000
Move	0.069	0.253	0.000	1.000
Sale conditional on move	0.780	0.414	0.000	1.000
No sale conditional on move	0.220	0.414	0.000	1.000
Sale conditional on move (t+1)	0.830	0.376	0.000	1.000
No sale conditional on move (t+1)	0.170	0.376	0.000	1.000
Became landlord conditional on move	0.121	0.326	0.000	1.000
Owner-occupied conditional on move	0.787	0.409	0.000	1.000
Rental conditional on move	0.122	0.327	0.000	1.000
Other tenure type conditional on move	0.091	0.287	0.000	1.000
Move within TWA	0.058	0.234	0.000	1.000
Move between TWAs	0.008	0.090	0.000	0.000
Move to undefined TWA	0.004	0.064	0.000	0.000
Move distance less than 50 km	0.057	0.231	0.000	1.000
Move distance over 50 km	0.008	0.088	0.000	0.000
End of unemployment	0.329	0.470	0.000	1.000
Move and end of unemployment	0.030	0.169	0.000	1.000
Move and end of unemployment, over 50km	0.005	0.069	0.000	0.000
Change of employer	0.149	0.356	0.000	1.000
Move and change employer	0.017	0.130	0.000	1.000
Move and change employer, over 50km	0.003	0.057	0.000	0.000

Notes: The table reports summary statistics for the key explanatory variables and outcomes for the full panel before bandwidth restrictions. The sample sizes for the outcome subsamples are as follows: unconditional (N = 1,752,331), conditional on move (N = 120,899), conditional on unemployment (N = 114,844) and conditional on employment (N = 1,190,167).

4 Empirical strategy

Our empirical strategy to analyze loss aversion follows a two-stage procedure common in the literature. The first stage is to estimate the predicted market value for each housing unit, and to calculate the expected rate of nominal return as the difference between the predicted market value and the initial purchase price. In the second stage, we analyze the relationship between expected return and mobility, and other outcomes of interest.

We improve on the accuracy of the first-stage prediction, compared to the previous literature, by using an ensemble machine learning method. We also provide transparent evidence of the performance of our pricing model. In the second stage, we examine visually mobility rates in bins of expected return and quantify the drop in mobility due to loss aversion through OLS-regression of mobility on a dummy for expected nominal loss while flexibly controlling for the rate of expected nominal return.

4.1 Predicting expected returns

We estimate the predicted market values for each apartment each year using an ensemble method that combines predictions from multiple individual models. Individual models are estimated using standard OLS, bootstrap aggregation of regression trees (bagging), random forest, and extreme gradient boosting (XGBoost).[11] Our motivation and the increasing interest in the economics discipline in machine learning and ensemble methods stems from their superior out-of-sample prediction performance compared to traditional linear models (e.g. Varian, 2014; Mullainathan and Spiess, 2017; Athey and Imbens, 2019).

We create separate ensemble models for each TWA. For each area we model the logarithm of the debt-free price of the apartment using detailed characteristics, the location of the apartment and year-fixed effects as explanatory variables. Before taking the logarithm, we deflate the prices to the base year (2005) to enhance the prediction performance. We convert the estimated housing prices back to nominal values before calculating the expected nominal returns. Previous literature has shown that the *nominal* purchase price is the key reference point for sellers.

In our model selection, we use a conventional random split data approach[12]. As the evaluation metric, we use root mean squared error (RMSE). In the data splitting, 80 percent of the transactions are assigned to the training sample that is used for model selection and for fitting the final models. The remaining 20 percent of transactions are assigned to the test sample, which is used only to provide a credible approximation of

[11]Our initial version also included elastic net but it was left out since it was not given any weight in the ensemble.

[12]In a machine learning context, model selection is commonly referred to as tuning. See e.g. Hastie et al. (2009).

the out-of-sample prediction performance. For OLS, we select the best-fitting hedonic specification using cross-validated RMSE. For the tree-based models, hyperparameters are tuned using out-of-bag error and Bayesian optimization. For computational efficiency, our k-fold cross-validation scheme uses 5 folds that are created using standard random assignment. The model selection procedure is performed separately for each TWA.

The ensemble is a weighted average of the predictions of individual methods. The weights are calculated using stacked regression (Breiman, 1996) as in Mullainathan and Spiess (2017) and Athey et al. (2019). In this model, the observed market price is regressed on stacked out-of-fold predictions from the selected individual methods, with weights constrained to be non-negative. van Der Laan et al. (2007) show that the prediction performance of the ensemble is asymptotically at least as good as the best individual prediction.

In earlier work, expected market prices have been retrieved through various methods. Survey responses would seem appealing as they directly represent homeowners' expectations. However, survey data often suffer from both small samples and measurement errors, which is also the case in Engelhardt (2003). Using cumulative changes in regional sales price indices (Steegmans and Hassink, 2018) requires less detailed data, but does not provide robust bases for prediction as this approach ignores apartment characteristics and local amenities that inherently determine the market value. The most common option is the hedonic approach, in which a market price is estimated using unit characteristics. Einiö et al. (2008) use a linear model with area and period fixed effects. Their paper, along with many others (Genesove and Mayer, 2001; Anenberg, 2011; Bokhari and Geltner, 2011), unfortunately lacks transparent evidence of the models' performance with unseen data.

4.2 Mobility analysis

Our identification strategy to isolate the impact of loss aversion from other determinants of mobility is based on the idea that the likelihood of an owner expecting to sell at a nominal loss increases sharply when our measure of expected rate of return is below zero. We argue that other determinants of mobility should develop smoothly around this cutoff. We start our empirical analysis with descriptive graphs that plot mobility rates and other outcomes in bins of expected returns. Then we estimate the magnitude of the discontinuity through OLS regression.

Our setting resembles a regression discontinuity design (Cattaneo and Titiunik, 2022; Imbens and Lemieux, 2008), but differs from an ideal RDD because treatment status is inaccurately measured. Despite the good predictive performance of our hedonic model, we may still erroneously classify apartments into the loss and gain domains close to the zero-gain cutoff, or owners may be uncertain as to whether to expect a loss or a gain.

Thus, a comparison of outcomes very close to the cutoff likely underestimates the effect of expecting a nominal loss. Because actual treatment status is not known, we cannot use the fuzzy RDD approach. To address the measurement issue, when estimating the model, we drop observations with expected rate of return between -5% and +5% in our base specification. This omitted region (or donut hole) is based on analysis of prediction errors in Section 5.1 (see footnote 15), and visual inspection of the binned outcomes. We analyze robustness to the size of the omitted region.

The model we estimate is written as follows:

$$Y_{it} = \alpha_t + \tau LOSS_{it} + \beta Return_{it} + \gamma(LOSS_{it} \times Return_{it}) + \mu X_{it} + \varepsilon_{it} \tag{1}$$

where Y_{it} is the mobility dummy or other outcome variable for household i in year t. $LOSS_{it}$ is a dummy variable that equals 1 if the expected rate of return is negative and τ is the treatment effect at the cutoff. $Return_{it}$ is the expected rate of nominal return, which we interact with the treatment indicator to allow for different slopes on either side of the cutoff. X_{it} includes controls for the indebtedness of the household through dummies for 5%-point bins of the loan-to-value ratio of the household and controls for the duration of the current residence spell. In addition, we control for year fixed effects. ε_{it} is the error term.[13] We report standard errors clustered at the municipality level (58 clusters).

Our main specification uses a bandwidth of -20% to +20% around the cutoff. This bandwidth is somewhat arbitrarily chosen, and we test for the robustness of our results to alternatives. Models developed for determining the optimal bandwidth in RDD settings are not well suited to our setting, where treatment status may be inaccurately measured in bins close to the cutoff, even after omitting the closest bins. Nevertheless, we also estimate the model with data-driven bandwidths to analyze the sensitivity of the findings to alternative bandwidths. In these checks we use bias-corrected estimates with robust inference, as suggested in Calonico et al. (2014). As it is unclear whether the effective sample size should be the same on both sides of the zero return cutoff, we test two different versions of MSE-optimal bandwidths for completeness.

Our design has similar potential threats to the validity of the analysis as standard RDDs. Firstly, there may be selection into the treatment group based on household characteristics related to mobility. Such selection could occur if different types of homeowners are differentially affected by loss aversion. Those who are strongly affected remain in the loss domain while others sell or move and exit the loss group. We address this issue by

[13]Earlier papers that study the mobility effects of nominal loss aversion use either linear probability or proportional hazard models. Engelhardt (2003) uses a simple linear probability model and a semi-parametric extended proportional hazards model. Steegmans and Hassink (2018) use also use extended (Cox) proportional hazard model. The hazard models require assumptions that expected returns and LTV are unlikely to satisfy. The linear model used requires less restrictive assumptions.

testing for the balance of covariates at the cutoff, and by adding household characteristics as controls.

Manipulation of the assignment variable by households would also bias the findings. One could hypothesize that those who consider selling adjust the quality through renovation, for example. These kind of quality changes are unobserved in the data and can lead to bias if they are more common on either side of the cutoff. We cannot rule out these kind of unobserved improvements, but it is not clear whether to expect a discontinuity in the propensity to renovate at the cutoff.

5 Results

5.1 Price model

Table 2 describes the predictive performance of our ensemble price model in unseen transaction data for each TWA.[14] In all areas, the absolute relative residual is below 20% for roughly 90 percent of observations and below 5% in around 40 percent of observations.

Table 2: Ensemble test sample predictive performance, 2006-2018

	Share of observations with absolute relative residual			
Area	< 0.05	< 0.10	< 0.15	< 0.20
Helsinki TWA	0.441	0.725	0.866	0.930
Tampere TWA	0.417	0.693	0.838	0.905
Turku TWA	0.363	0.645	0.803	0.885

Notes: The table reports the share of test sample observations where the ensemble model's absolute relative residual, defined as $|1-\exp(\hat{p}-p)|$, is below each threshold.

The ultimate purpose of the prediction model is to produce an estimate of the expected rate of nominal return for each household in each year in the household data. We cannot directly assess the performance of the model in this respect as households' perceptions are unobservable. However, statistics on prediction errors in the transaction data in Table 2 can be used to infer the likelihood that the true expected return has a sign opposite to our prediction in different parts of the distribution of predicted returns. The table indicates that the likelihood of a wrong sign decreases steeply close to the cutoff, and more slowly further away from the cutoff as the likelihood of a wrong sign approaches zero.[15] This means that we can plausibly assume that there is an abnormally sharp

[14] Table A3 also shows the results for all the individual models.

[15] For example in Helsinki, 56% (100%-44%) of absolute relative prediction errors exceed 5% in the transaction data. Extrapolating to the household data and assuming symmetric prediction errors, this

increase in the probability of expecting a nominal loss when we move from a positive to a negative predicted return. Hence the estimation method discussed in Section 4 should work.

Table A5 reports the ensemble weights assigned to individual models in each TWA. The largest weights are assigned to the decision tree-based estimation methods XGBoost and random forest, which was expected based on earlier machine learning literature. Only in the Tampere TWA, a minor non-zero weight of 1.2% is assigned to OLS.

The test set RMSE in Table A4 for the ensemble indicates that on average our model results in a 12.4% prediction error in apartment values for unseen data in the Helsinki TWA. The improvement over traditional OLS is 22%. For the Tampere and Turku TWAs, this improvement is 17% and 23% respectively. These findings show that an ensemble of tree-based models is superior in a setting where accurate out-of-sample prediction is important for the credibility of the analysis.

5.2 Predicted returns

Figure 2 shows the distribution of realized and predicted nominal returns for the transactions in the matched apartment-household panel. The realized returns show bunching at zero gain and slightly above zero, and missing density left of zero. The predicted returns line illustrates the distribution of returns as if households sold at the predicted market value, and it does not result in similar bunching. This can be interpreted as first evidence of loss aversion.[16]

Figure 3 shows the distribution of expected return in the matched apartment-household data, regardless of whether the apartment was sold. The shape of the distribution is similar to the distribution for the realized transactions in Figure 2. Even if the general price trend in all TWAs was positive in the period studied, a substantial share of households can expect to sell at a loss. This could occur, for example, if they bought an apartment with characteristics that have become less scarce in the TWA, or in an area within the TWA where prices have declined. In addition, deterioration of the condition of the apartment through aging can lead to a decrease in expected price even if the general trend is up.[17] The density discontinuity test does not reject the null hypothesis of continuity at the zero returns cutoff with Cattaneo et al. (2020) optimal bandwidths.[18]

would mean that 28% (half of 56%) of households with a predicted return of -5% actually expect a positive return. For households with a -10% predicted return, the share with a positive expected return would be 13,75% (half of 100%-72.5%). For households with a -15% predicted return, the share of positive expected return would be 6,7%.

[16]The evidence is more pronounced in Figure A2, which shows the distribution of the difference of realized and predicted returns.

[17]Figure A1 in the Appendix shows the mass of households with predicted loss and gain by TWA and year.

[18]With a standard McCrary (2008) test the null hypothesis of continuity is rejected (Z-statistic = 4.48).

Figure 2: Bunching in nominal capital return distribution

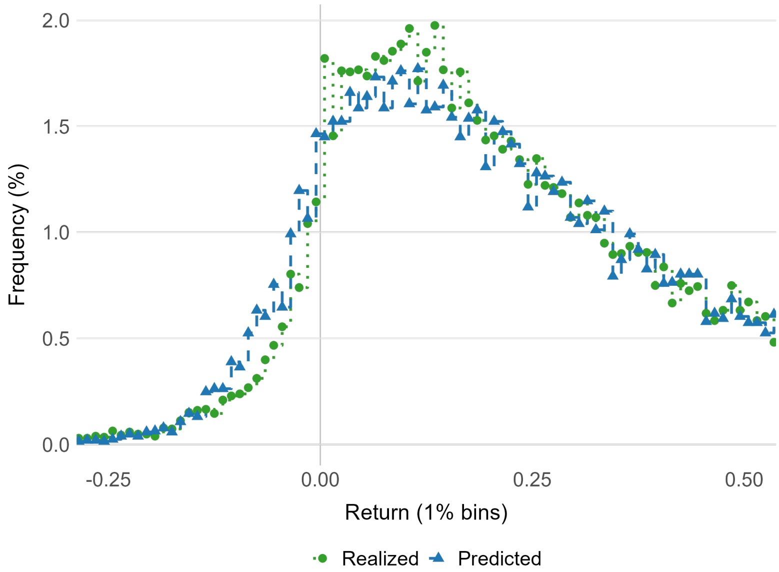

Notes: Figure plots the distribution of realized (dotted green) and predicted (dashed blue) returns in 1% bins for transactions in the matched household-apartment panel. Zero return marked by vertical line. Transactions included in the training sets of the ensemble model are excluded from the figure.

Figure 3: Continuity of predicted nominal capital gains in the apartment-household data

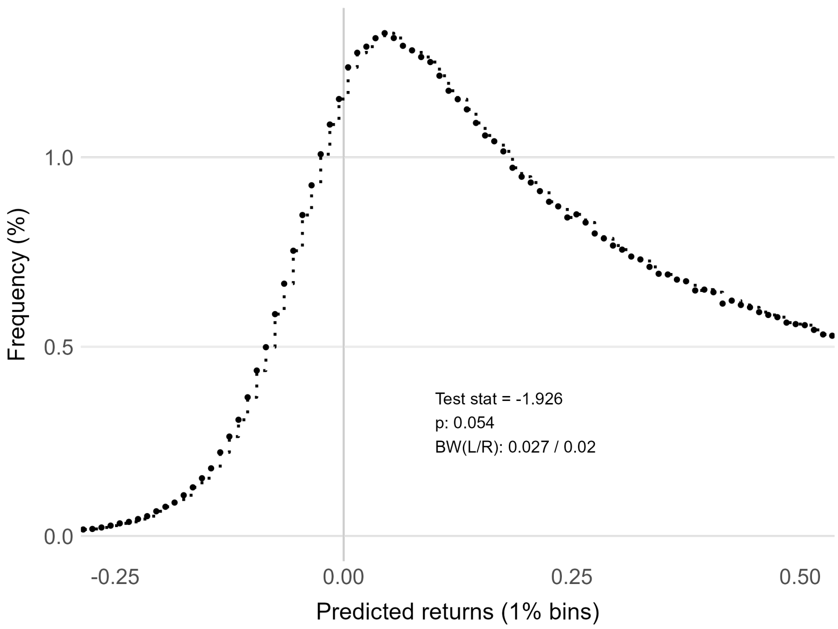

Notes: Figure shows the distribution of homeowners' predicted nominal capital returns around the zero threshold for 2006–2018, along with the density discontinuity test results following the method of Cattaneo et al. (2020).

5.3 Mobility

We start our main analysis with visual examination of discontinuities around the zero return cutoff. We plot the coefficients on dummies for expected return bins (2%) for different outcomes while controlling for LTV, spell length, and year. We are interested in discontinuities at the cutoff. The level of the coefficients is not informative as it is affected by the controls.

The upper panels of Figure 4 plot the coefficients on dummies for expected return bins for sale and move responses. The figure shows a very similar pattern for both outcomes, with a clear jump around the zero-return cutoff, but this jump appears to be smaller for mobility than for sales. Panel (c) shows a similar graph for the sale decision (in year t or t+1), but now conditional on moving in year t. The likelihood of selling when moving is lower below the threshold, which indicates that some households willing to move but predicted to sell at a loss leave their previous apartment unsold to avoid realizing a nominal loss.

Figure 4: Sale and move

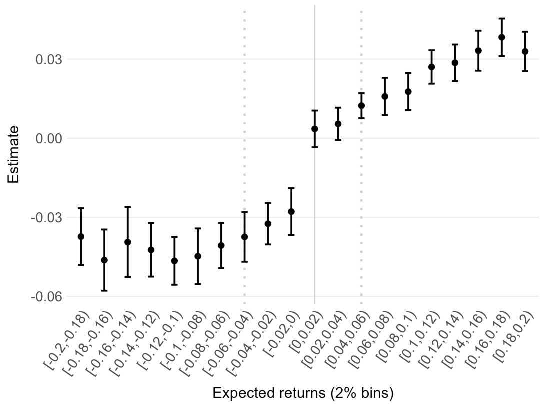

(a) Sale

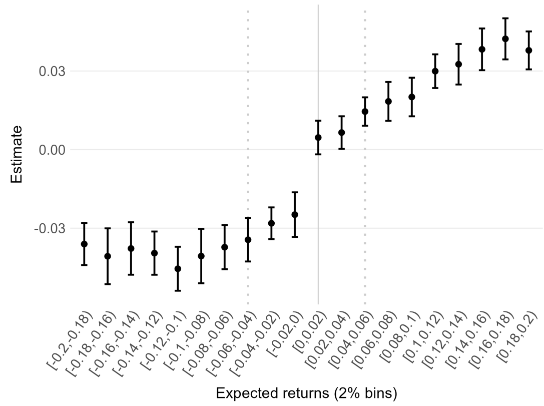

(b) Move

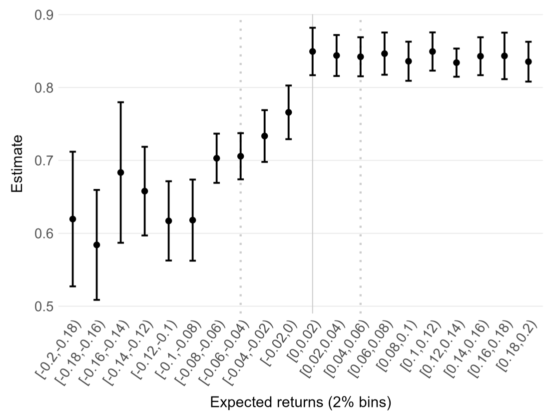

(c) Sale conditional on move, (t+1)

Notes: Figure shows estimates ($\hat{\beta}_i$) for expected return bins (2% intervals) from regressions of the form $Y_{it} = \sum_i \beta_i \, \mathrm{Bin}_i + \gamma \, X_{it} + \delta \, \mathrm{Year} + \varepsilon_{it}$, where X include controls for spell length (entered in linear and squared form) and loan-to-value ratio (in 5% bins). Solid vertical lines indicate the bin containing zero expected return. Dotted vertical lines indicate $\pm$ 5% expected return, which is used as an omitted (donut) region in our main specification. Error bars show 95% confidence intervals based on standard errors clustered at the municipality level.

Figure 5 further explores the differences in the housing decisions of movers around the cutoff. In panel (a), we see that movers with a predicted loss are more likely to start receiving rental income after moving. This suggests that renting out the previous apartment instead of selling is one way to avoid realizing nominal losses when moving. Panels (b) and (c) are less clear, but indicate that movers with a predicted loss may be less likely to continue as homeowners and more likely to become renters. This could be because loss-averse movers who do not sell their previous apartment are unable to buy a new apartment due to credit constraints.

Figure 5: Tenure type of previous and new housing unit

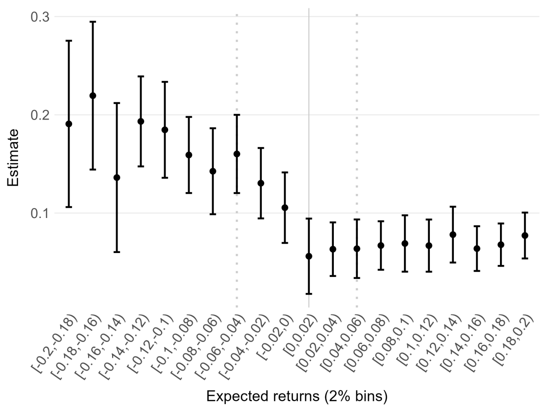

(a) Became landlord conditional on move

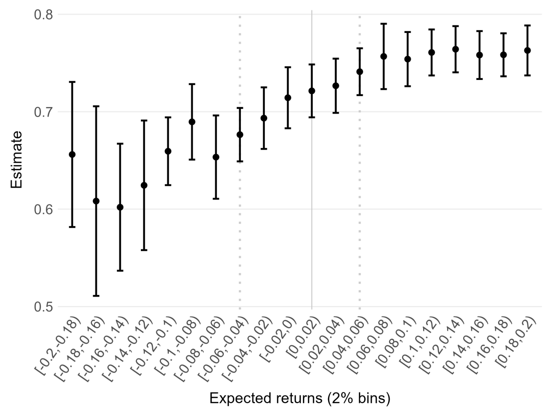

(b) Owner-occupied conditional on move

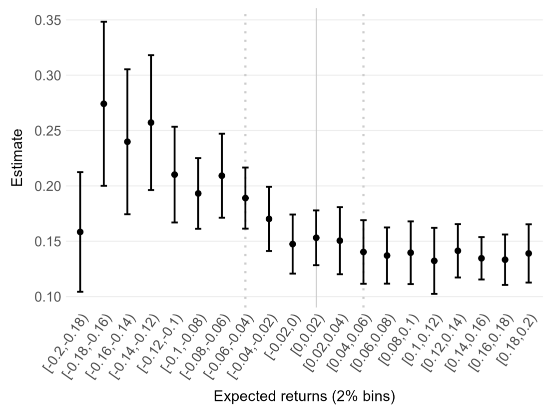

(c) Rental conditional on move

Notes: Figure shows estimates ($\hat{\beta}_i$) for expected return bins (2% intervals) from regressions of the form $Y_{it} = \sum_i \beta_i \, \mathrm{Bin}_i + \gamma X_{it} + \delta \, \mathrm{Year} + \varepsilon_{it}$, where X include controls for spell length (entered in linear and squared form) and loan-to-value ratio (in 5% bins). Solid vertical lines indicate the bin containing zero expected return. Dotted vertical lines indicate ± 5% expected return, which is used as an omitted (donut) region in our main specification. Error bars show 95% confidence intervals based on standard errors clustered at the municipality level.

Panels (a) and (b) of Figure 6 analyze mobility within and between TWAs. Both types of mobility seem to be affected by loss aversion. The results are very similar when we divide moves to below 50 km moves and over 50 km moves in the lower panel. Moves within region or short distance moves are likely to be mainly motivated by adjustment of housing consumption, e.g. upsizing or downsizing. Thus, a reduction in moves within the region or short distance moves arguably leads to mismatch in the housing market. The fact that we also find an effect on moves between TWAs and long distance moves suggests that loss aversion may affect the functioning of the labor market by locking in workers in their current labor market areas.

Figure 7 provides more direct evidence on labor market implications of loss aversion. Panels (a) and (b) show no clear effects on the propensity to change employers or the end of unemployment, although there is some indication of a small reduction in job changes. Nevertheless, loss aversion does not seem to be an important driver of these labor market transitions. It should, however, be noted that the bulk of job changes and unemployment endings do not coincide with moves, which means that effects on the small share of labor market transitions linked to moves may not be detectable when analyzing all job changes and unemployment endings. Panels (c) and (d) analyze moves linked with labor market transitions, and we find negative effects on moving and changing employer and moving and becoming employed. Whether these graphs can be interpreted as evidence of labor market effects is questionable, because most moves are of short distance, and the labor market transition linked with short moves could have taken place even without moving. Panels (e) and (f) address this issue by focusing on long distance moves (over 50 km) combined with labor market transitions. Panel (e) shows that longer distance moves to change employer are less common in the loss domain, but there is no clear effect on long distance moves linked with end of unemployment.

Table 3 reports the main regression results for the specification (1).[19] The purpose of the table is to quantify the discontinuities observed in binned plots of outcomes of interest. The estimation relies on extrapolating the fit of the regression lines from data outside the omitted region between -5% and +5% to the omitted region and assumes a piecewise linear relationship between expected returns and the outcome variable. Thus, the model departs from an ideal RDD and the estimates should be interpreted cautiously. Section 5.5 analyzes the sensitivity of the results to various modeling choices.

The point estimate in the first row of Table 3 suggests that expecting a loss decreases homeowners' propensity to sell by 4.1 percentage points. This corresponds to a 51% decrease in the mean selling propensity compared to the control mean. The reduction in mobility rate due to loss-exposure is somewhat smaller (4.0 percentage points or 43.5%). Conditional on moving, loss-exposure reduces the propensity to sell by 11.5 percentage points or 14%. The effect on the propensity to become a landlord conditional on moving is substantial: 7.1 percentage points or 61%. Table A10 in the appendix analyzes the persistence of the effect on becoming landlord and no sale conditional on moving. The effect on the propensity to become landlord first increases in t+2 but drops in t+3. The

[19]Table A11 in the appendix shows results from the same specification, but with predicted returns based on standard OLS.

Figure 6: Mobility within and across regions

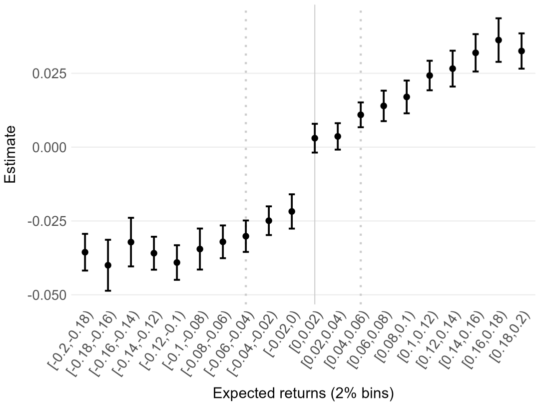

(a) Move within TWA

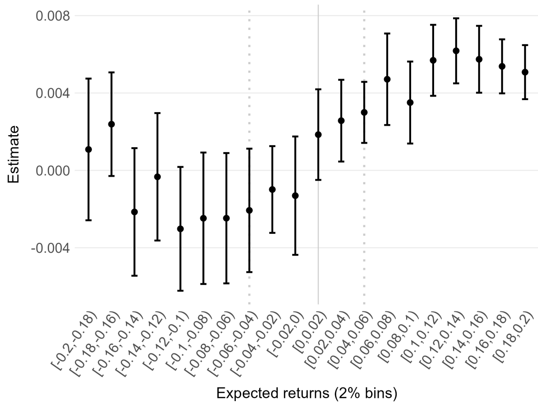

(b) Move between TWAs

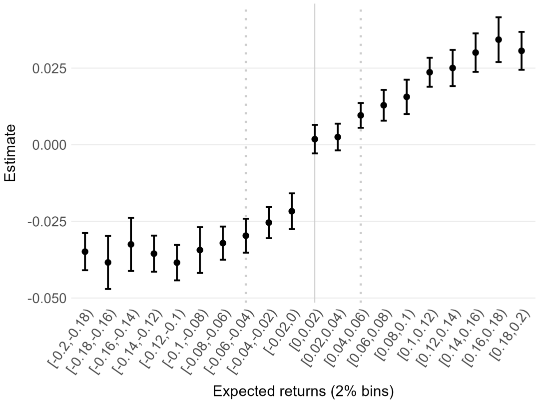

(c) Move below 50 km

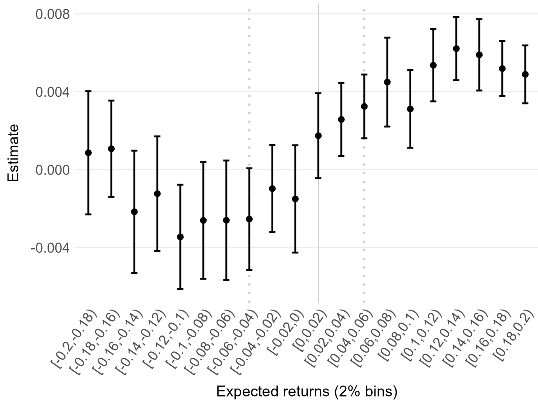

(d) Move over 50 km

Notes: Figure shows estimates ($\hat{\beta}_i$) for expected return bins (2% intervals) from regressions of the form $Y_{it} = \sum_i \beta_i \, \text{Bin}_i + \gamma X_{it} + \delta \, \text{Year} + \varepsilon_{it}$, where X include controls for spell length (entered in linear and squared form) and loan-to-value ratio (in 5% bins). Solid vertical lines indicate the bin containing zero expected return. Dotted vertical lines indicate ± 5% expected return, which is used as an omitted (donut) region in our main specification. Error bars show 95% confidence intervals based on standard errors clustered at the municipality level.

effect on not selling conditional on move weakens gradually and roughly halves by t+3.[20] When we estimate the model for tenure status of the new dwelling, we find that expected loss increases the propensity to become a renter conditional on moving by 4.3 percentage points (or 34%).

Next we analyze separately the effects on within TWA moves that are likely to be motivated by adjustments of housing consumption, and moves between TWAs that are more likely to be motivated by labor market considerations. Both within and between TWAs moves are strongly affected by loss-exposure. The proportionate impact relative to the baseline is stronger for moves between TWAs. The findings are very similar when we divide moves based on the distance of move (over or below 50 km). These findings

[20] Note that the sample used in Table A10 is smaller than in Table 3 because we keep only observations that can be followed three years after move in Table A10.

Figure 7: Labor market transitions

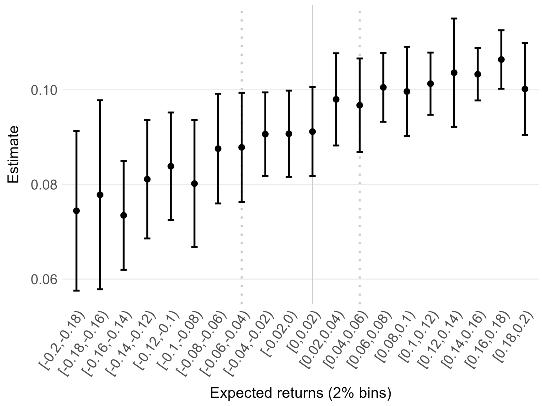

(a) Change of employer

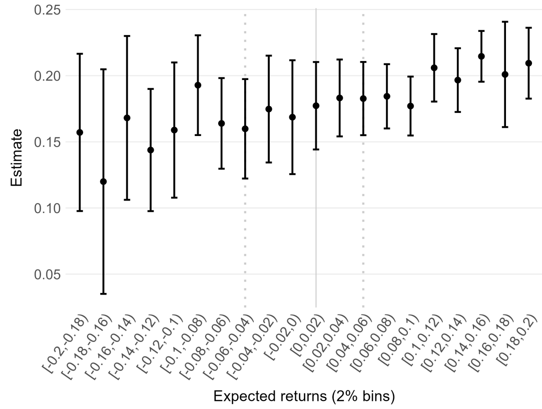

(b) End of unemployment

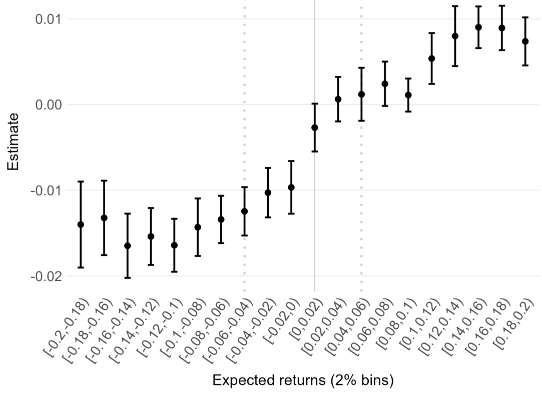

(c) Move and change employer

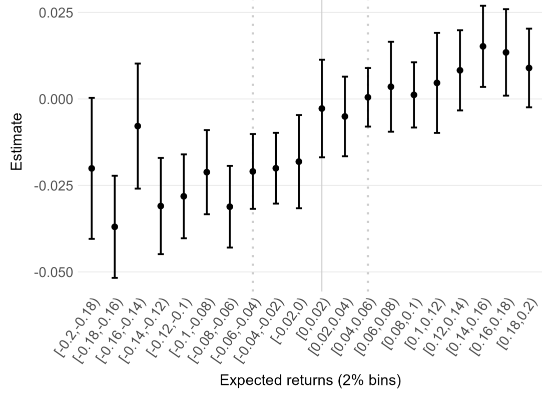

(d) Move and end of unemployment

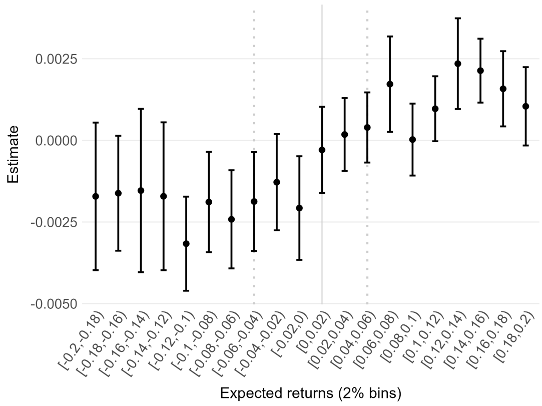

(e) Move over 50km and change employer

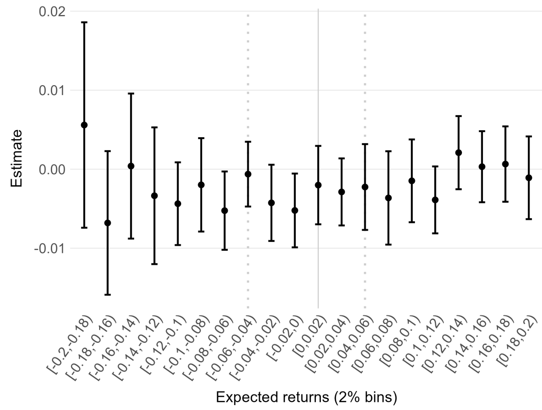

(f) Move over 50km and end of unemployment

Notes: Figure shows estimates ($\hat{\beta}_i$) for expected return bins (2% intervals) from regressions of the form $Y_{it} = \sum_i \beta_i \, \text{Bin}_i + \gamma \, X_{it} + \delta \, \text{Year} + \varepsilon_{it}$, where X include controls for spell length (entered in linear and squared form) and loan-to-value ratio (in 5% bins). Solid vertical lines indicate the bin containing zero expected return. Dotted vertical lines indicate $\pm$ 5% expected return, which is used as an omitted (donut) region in our main specification. Error bars show 95% confidence intervals based on standard errors clustered at the municipality level.

indicate that loss aversion may have implications for the functioning of the labor market by impeding relocation of the labor force.

To get more direct evidence of labor market effects, we analyze two important labor market transitions: end of unemployment and change of employer. The impact on the overall propensity to become employed when unemployed or the propensity to change employer is insignificant, but this may be because only a small fraction of these transitions involve a move. The impact on moves where unemployment ends at the same time is sizable. However, when the move was of short distance, the unemployment could have ended even without moving, because the new job was likely as accessible both before and after the move. Therefore we next analyze longer distance moves (over 50 km) where unemployment ended, and find no effect. This suggests that loss aversion does not prevent unemployed homeowners from accepting jobs further away from their current location. For the propensity to move and change employer, we find a negative effect that persists when we omit short distance moves. This finding indicates that loss aversion may contribute to regional mismatch in the labor market by discouraging homeowners from taking up new (potentially more productive) jobs that require moving. Long distance moves linked with job changes are, however, rare among homeowners, and the effect of loss-exposure is economically rather small, only 0.3 percentage points.

Table 3: Main regression results

Outcome	Obs	Control mean	Loss
Sale	386,698	0.080	-0.041 *** (0.003)
Move	386,698	0.092	-0.040 *** (0.003)
Sale conditional on move (t+1)	34,076	0.839	-0.115 *** (0.020)
No sale conditional on move (t+1)	34,076	0.161	0.115 *** (0.020)
Became landlord conditional on move	34,076	0.116	0.071 *** (0.020)
Owner-occupied conditional on move	34,076	0.778	-0.043 ** (0.020)
Rental conditional on move	34,076	0.126	0.043 ** (0.017)
Move within TWA	386,698	0.076	-0.030 *** (0.002)
Move between TWAs	386,698	0.012	-0.008 *** (0.001)
Move distance less than 50 km	386,698	0.074	-0.029 *** (0.002)
Move distance over 50 km	386,698	0.011	-0.008 *** (0.001)
End of unemployment	26,185	0.413	0.023 (0.025)
Move and end of unemployment	26,185	0.045	-0.021 ** (0.008)
Move and end of unemployment, over 50km	26,185	0.006	0.001 (0.003)
Change of employer	304,703	0.172	-0.008 (0.006)
Move and change employer	304,703	0.022	-0.010 *** (0.002)
Move and change employer, over 50km	304,703	0.005	-0.003 *** (0.001)

Notes: The table reports estimates for τ from equation (1): $Y_{it} = \alpha_t + \tau LOSS_{it} + \beta Return_{it} + \gamma(LOSS_{it} \times Return_{it}) + \mu X_{it} + \varepsilon_{it}$. $LOSS_{it}$ and $Return_{it}$ use fitted values from a hedonic price model estimated by the ensemble model. The estimation uses a bandwidth of -20% to 20% around the zero return cutoff, excluding observations in the donut hole $[-5\%, 5\%]$. The control mean corresponds to the predicted outcome just above the zero return cutoff with control variables X_{it} set to their sample means. Estimates for β and γ are reported separately in Table A9. Standard errors clustered at the municipality level are in parentheses. Significance is denoted by asterisks: *p<.1,**p<.05 ,***p<.01.

5.4 Heterogeneity

Next we analyze the heterogeneity of sale and mobility responses with respect to household leverage and length of ownership spell. High leverage could reinforce the effect of loss aversion through a tighter credit constraint when buying a new apartment. Longer residence spell means that the initial purchase price is more distant and could thereby diminish its importance as a reference point for the seller. In addition, it is of interest to analyze heterogeneity of responses related to housing choices conditional on moving.

Figure 8 plots the coefficients for the dummy variables corresponding to the expected return bins (2%) related to mobility, now estimated for subsamples classified by the LTV ratio and the duration of the ownership spell. Visual inspection suggests that high LTV ($> 50\%$) and short ownership spell (< 5 years) correspond to a clearer jump around the threshold (panel (c)). Figure A4 shows similar patterns for sale. As regards housing choices, Figures A5 and A6 suggest that high LTV strengthens the impact of expected loss on the likelihood of no sale conditional on move and the likelihood of becoming a landlord conditional on move. For the propensity to become a renter conditional on move (Figure A7) we do not find clear differences across sub-samples.

Table 4 quantifies the heterogeneous effects by interacting the dummy for expected loss with dummies for high LTV and long spell. For both sale and move decisions, the impact of a nominal loss is clearly more strongly negative (4.6 percentage points) with high LTV. A long spell leads to an equally large decrease in the effect of expected loss. High LTV increases the effect of expected loss on the propensity to become a landlord by 8 percentage points, and the effect on the propensity to not sell conditional on moving by almost as much. A plausible explanation could be that homeowners with high LTV are especially reluctant to sell at a loss, because they may not be able to pay back their mortgage, and rent out their previous apartment instead. Interestingly, a high LTV does not significantly affect the impact of loss-exposure on the propensity to become a renter conditional on move.

Table 4: Heterogeneity: Spell length and LTV

Outcome	Control mean	Loss	Loss X LTV>0.50	Loss X Spell length>5
Sale	0.079	-0.022 *** (0.003)	-0.046 *** (0.002)	0.051 *** (0.002)
Move	0.092	-0.022 *** (0.004)	-0.046 *** (0.002)	0.055 *** (0.002)
Became landlord conditional on move	0.118	0.025 (0.021)	0.081 *** (0.009)	-0.059 *** (0.011)
No sale conditional on move (t+1)	0.163	0.077 ** (0.029)	0.068 *** (0.025)	-0.057 *** (0.017)
Rental conditional on move	0.125	0.042 ** (0.019)	-0.009 (0.012)	0.038 ** (0.016)

Notes: The table reports estimates for τ, ς and η, from equation $Y_{it} = \alpha_t + \tau LOSS_{it} + \beta Return_{it} + \gamma(LOSS_{it} \times Return_{it}) + \varsigma(LOSS_{it} \times D(LTV_{it} > 0.5)) + \eta(LOSS_{it} \times D(Spell_{it} > 5)) + \mu X_{it} + \varepsilon_{it}$. The estimation uses a bandwidth of -20% to 20% around the zero return cutoff, excluding observations in the donut hole $[-5\%, 5\%]$. The control mean corresponds to the predicted outcome just above the zero return cutoff with control variables X_{it} set to their sample means. Standard errors clustered at the municipality level are in parentheses. Significance is denoted by asterisks: *p<.1,**p<.05 ,***p<.01.

5.5 Validity and robustness checks

Next we test whether the socio-economic characteristics of the household and housing characteristics develop smoothly at the cutoff. Figure A8 shows binned plots of the mean age of the owners, dummy for children, dummy for higher education, dummy for shared ownership of the apartment, as well as floor area of the apartment and age of the building.

Figure 8: Heterogeneous mobility responses by LTV and spell length

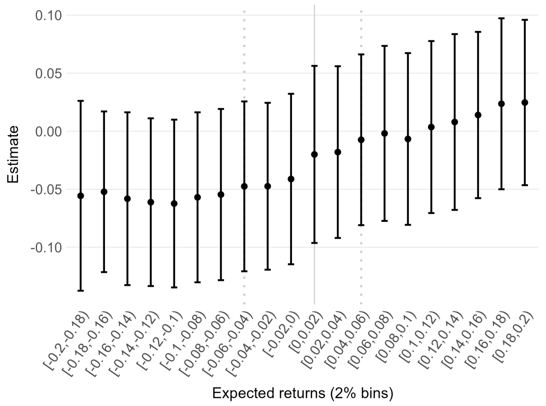

(a) Low LTV, Short Spell

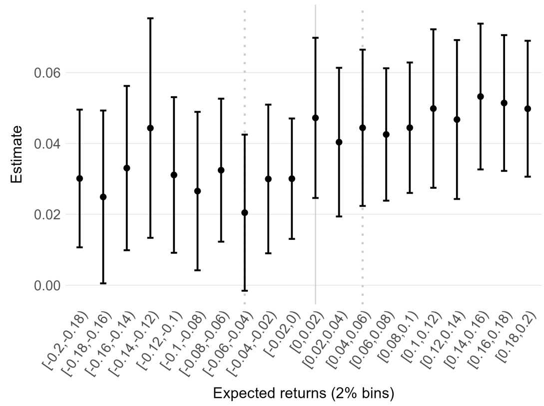

(b) Low LTV, Long spell

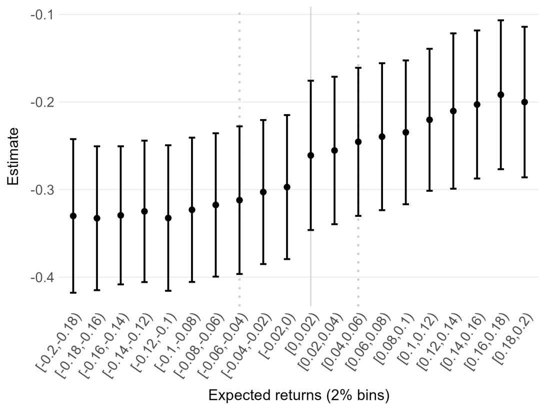

(c) High LTV, Short spell

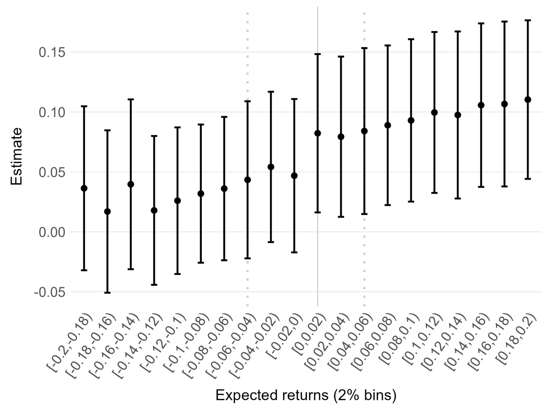

(d) High LTV, Long spell

Notes: Figure shows estimates ($\hat{\beta}_i$) for expected return bins (2% intervals) from regressions of the form $Y_{it} = \sum_i \beta_i \,\mathrm{Bin}_i + \gamma\, X + \delta\, \mathrm{Year} + \varepsilon$, fitted to sub-samples based on LTV and spell length. X include controls for spell length (entered in linear and squared form) and loan-to-value ratio (in 5% bins). Solid vertical lines indicate the bin containing zero expected return. Dotted vertical lines indicate $\pm$ 5% expected return, which is used as an omitted (donut) region in our main specification. Error bars show 95% confidence intervals based on standard errors clustered at the municipality level.

Table 5 shows the coefficients of the loss dummy for these variables from eq (1). Most covariates are balanced, but we find statistically significant discontinuities for owner age and floor area of the apartment. This suggests that there may be some selection on household and apartment characteristics into the loss group. However, these discontinuities are quantitatively small and in Table A11 we find that adding these characteristics or spell length dummies as additional controls does not change the results. This alleviates the concern that selection is an important driver of our findings.

Table 5: Covariate balance

Outcome	Obs	Loss
Avg. owner age	386,698	0.895 *** (0.224)
Children	386,698	0.008 (0.009)
Higher education	386,698	-0.002 (0.011)
2 owners	386,698	0.014 * (0.008)
Household disposable income	386,698	-576.042 (602.417)
Floor area (sqm)	386,698	2.332 *** (0.356)
Building age	386,698	0.212 (0.549)

Notes: The table reports estimates for τ, from equation $Y_{it} = \alpha_t + \tau LOSS_{it} + \beta Return_{it} + \gamma(LOSS_{it} \times Return_{it}) + \varepsilon_{it}$. The estimation uses a bandwidth of -20% to 20% around the zero return cutoff, excluding observations in the donut hole $[-5\%, 5\%]$. Standard errors clustered at the municipality level are in parentheses. Significance is denoted by asterisks: *p<.1,**p<.05 ,***p<.01.

Estimates in Table 3 assume a linear relationship between the predicted gain and the outcome variables. The slope is allowed to differ on either side of the cutoff. This assumption is relaxed in Table A11 (last column) that reports estimates where we add the second-order term of the expected loss (and its interaction with the $LOSS$ dummy). The estimates are of the same sign as with the linear specification, but the proportionate effects of expected loss become weaker across the board. These findings show that our quantitative results are somewhat sensitive to functional form. However, the second-order polynomial specification may be problematic in our setting where we measure loss status with error and the likelihood of error is closer to the cutoff. Larger share of erroneously predicted loss status close to the cutoff could bend the second-order polynomials on the left and right of the cutoff towards each other, and bias the estimates of the impact of loss-exposure towards zero.

Our main results use a somewhat arbitrarily chosen bandwidth of -20 to 20 percentage points around the zero gain cutoff, and an omitted region of -5 to 5 percentage points. Figure 9 shows the robustness of our results to alternative bandwidths and omitted regions around the cutoff for four key outcomes. Reducing the bandwidth from 20 percentage points to 15 percentage points does not affect the point estimates but increases the standard errors. Using a narrower donut hole leads to slightly smaller point estimates. In specifications without a donut hole, our estimates are close to zero. This is expected as close to the cutoff the treatment is measured with error attenuating the estimates. Using the robust RDD method by Calonico et al. (2014) (mserd and msetwo) leads to qualitatively similar but less precise results, especially when we use the 5 percentage point donut hole. Figures A9, A10 and A11 report similar robustness checks for other outcomes.

Figure 9: Sale and move

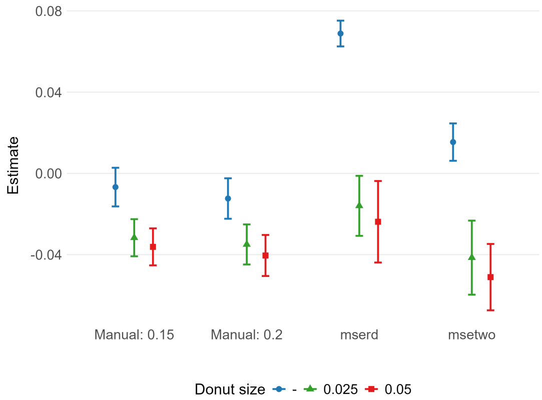

(a) Sale

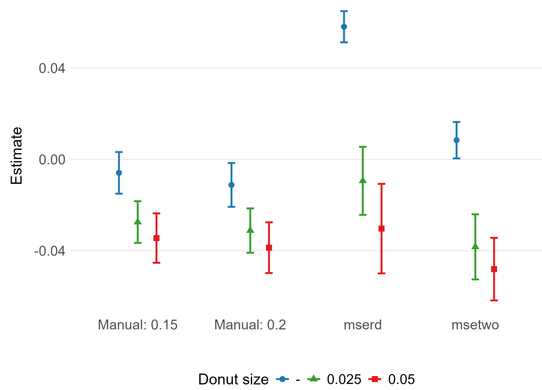

(b) Move

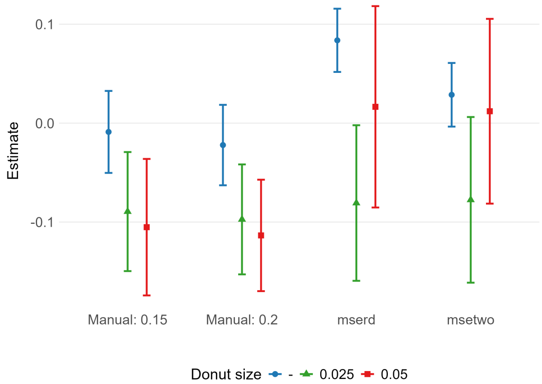

(c) Sale conditional on move, (t+1)

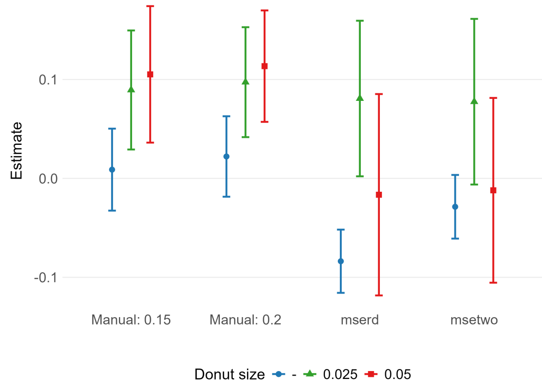

(d) No sale conditional on move, (t+1)

Notes: The figure shows local polynomial estimates of gains with conventional CIs for manual bandwidth models and bias-corrected estimates of gains with robust CIs for MSE-optimal models (Calonico et al., 2014). "mserd" uses symmetric bandwidths, while "msetwo" allows asymmetric bandwidths. All models use first-order polynomials, uniform kernel weighting, and control for LTV dummies, spell length, and year dummies.

6 Conclusions

We provide new evidence on the effect of nominal loss aversion on household mobility. We find that loss aversion affects homeowners' propensity to sell more strongly than the propensity to move. Both effects are economically important. Renting out the previous apartment seems to allow homeowners to avoid realizing nominal losses when moving. There is also some indication that homeowners with an expected loss are more likely to become renters conditional on moving. Our findings imply that loss aversion leads to misallocation of housing by locking homeowners with an expected loss into their current apartments. However, our results suggest that a well-functioning private rental market can mitigate this effect by offering loss-averse homeowners opportunities to move without selling their previous unit. At the same time, the composition of their investment portfolios is tilted towards rental housing. Increased flow to rental housing in the loss domain suggests that loss aversion can shape tenure type distribution in the housing market.

Heterogeneity analysis shows that the mobility effects of loss aversion are stronger among homeowners with high LTV and short ownership spells. These findings point to the interplay between credit constraints and loss aversion, and the reduced importance of the initial purchase price as a reference point over time.

An expected loss reduces not only intra-regional but also inter-regional mobility, which suggests that loss aversion can lead to misallocation of the labor force. We do not find direct evidence of loss aversion affecting flows out of unemployment through inter regional moves, but loss-exposure reduces job changes associated with inter-regional moves.

Accurate prediction of sales prices for potential sales is very important for analyzing loss aversion. Previous studies have used standard OLS hedonic models or regional price indices for prediction. We show that ensemble machine learning methods vastly outperform these methods.

We find that even in a period with rising nominal prices, a significant share of homeowner households can be expected to sell at a loss. Thus, nominal loss aversion is not confined to housing busts, but also affects the functioning of the market in normal times. A back-of-the-envelope calculation[21] using our preferred estimates for the mobility effect implies that in our context the overall mobility rate of homeowners would be almost 6% higher in the absence of loss aversion. During a housing market bust, the impact on the overall rate of mobility is likely to be an order of magnitude larger.

[21] Mean annual mobility rate in our main sample is 6.9% and 10.4% are predicted to sell at loss. Applying our estimate of 4 percentage point reduction in mobility due to loss aversion suggests that without loss aversion, overall mobility rate would be 0.4 percentage points higher, which corresponds to a 5.8% increase in overall mobility.

References

Andersen, S., Badarinza, C., Liu, L., Marx, J., Ramadorai, T., 2022. Reference dependence in the housing market. American Economic Review 112, 3398–3440.

Anenberg, E., 2011. Loss aversion, equity constraints and seller behavior in the real estate market. Regional Science and Urban Economics 41, 67–76.

Athey, S., Bayati, M., Imbens, G., Qu, Z., 2019. Ensemble methods for causal effects in panel data settings. AEA Papers and Proceedings 109, 65–70.

Athey, S., Imbens, G.W., 2019. Machine learning methods that economists should know about. Annual Review of Economics 11, 685–725.

Banki, D., Simonsohn, U., Walatka, R., Wu, G., 2025. Decisions under risk are decisions under complexity: Comment. Available at SSRN 5127515 .

Bokhari, S., Geltner, D., 2011. Loss aversion and anchoring in commercial real estate pricing: Empirical evidence and price index implications. Real Estate Economics 39, 635–670.

Bracke, P., Tenreyro, S., 2021. History dependence in the housing market. American Economic Journal: Macroeconomics 13, 420–43.

Breiman, L., 1996. Stacked regressions. Machine learning 24, 49–64.

Calonico, S., Cattaneo, M.D., Titiunik, R., 2014. Robust nonparametric confidence intervals for regression-discontinuity designs. Econometrica 82, 2295–2326.

Cattaneo, M.D., Jansson, M., Ma, X., 2020. Simple local polynomial density estimators. Journal of the American Statistical Association 115, 1449–1455.

Cattaneo, M.D., Titiunik, R., 2022. Regression discontinuity designs. Annual Review of Economics 14, 821–851.

van Der Laan, M.J., Polley, E.C., Hubbard, A.E., 2007. Super learner. Statistical Applications in Genetics and Molecular Biology 6.

Eerola, E., Harjunen, O., Lyytikäinen, T., Saarimaa, T., 2021. Revisiting the effects of housing transfer taxes. Journal of Urban Economics 124, 103367.

Einiö, M., Kaustia, M., Puttonen, V., 2008. Price setting and the reluctance to realize losses in apartment markets. Journal of Economic Psychology 29, 19–34.

Engelhardt, G.V., 2003. Nominal loss aversion, housing equity constraints, and household mobility: evidence from the united states. Journal of Urban Economics 53, 171–195.

Ferreira, F., Gyourko, J., Tracy, J., 2010. Housing busts and household mobility. Journal of Urban Economics 68, 34–45.

Fonseca, J., Liu, L., 2024. Mortgage lock-in, mobility, and labor reallocation. The Journal of Finance 79, 3729–3772.

Genesove, D., Mayer, C., 2001. Loss aversion and seller behavior: Evidence from the housing market. The Quarterly Journal of Economics 116, 1233–1260.

Goodman, L.S., Mayer, C., 2018. Homeownership and the american dream. Journal of Economic Perspectives 32, 31–58.

Greenaway-McGrevy, R., Sorensen, K., 2021. A time-varying hedonic approach to quantifying the effects of loss aversion on house prices. Economic Modelling 99, 105491.

Hastie, T., Tibshirani, R., Friedman, J.H., 2009. The elements of statistical learning: data mining, inference, and prediction. volume 2. Springer.

Hilber, C.A., Lyytikäinen, T., 2017. Transfer taxes and household mobility: Distortion on the housing or labor market? Journal of Urban Economics 101, 57–73.

Ihlanfeldt, K.R., 2011. Do caps on increases in assessed values create a lock-in effect? evidence from florida's amendment one. National Tax Journal 64, 7–25.

Imbens, G.W., Lemieux, T., 2008. Regression discontinuity designs: A guide to practice. Journal of econometrics 142, 615–635.

Kahneman, D., Tversky, A., 1979. Prospect theory: An analysis of decision under risk. Econometrica 47, 263–292.

McCrary, J., 2008. Manipulation of the running variable in the regression discontinuity design: A density test. Journal of econometrics 142, 698–714.

Mullainathan, S., Spiess, J., 2017. Machine learning: An applied econometric approach. Journal of Economic Perspectives 31, 87–106.

Oprea, R., 2024. Decisions under risk are decisions under complexity. American Economic Review 114, 3789–3811.

Oswald, A.J., 1996. A conjecture on the explanation for high unemployment in the industrialized nations: Part i. Warwick Economic Research Papers 475.

Quigley, J.M., 1987. Interest rate variations, mortgage prepayments and household mobility. The Review of Economics and Statistics , 636–643.

Ross, S.L., Zhou, T., 2021. Loss Aversion in Housing Sales Prices: Evidence from Focal Point Bias. Technical Report. National Bureau of Economic Research.

Steegmans, J., Hassink, W., 2018. Decreasing house prices and household mobility: An empirical study on loss aversion and negative equity. Journal of Regional Science 58, 611–634.

Stein, J.C., 1995. Prices and trading volume in the housing market: A model with down-payment effects. The Quarterly Journal of Economics 110, 379–406.

Thaler, R.H., 2016. Behavioral economics: Past, present, and future. American economic review 106, 1577–1600.

Varian, H.R., 2014. Big data: New tricks for econometrics. Journal of Economic Perspectives 28, 3–28.

Appendix

A1.1 Data

A1.1.1 Sources

The data sources are described below in Table A1. The main data set is the transfer tax register for 2006-2018 provided by the Finnish Tax Administration[22]. The register provides reliable information on all market transactions and ownership spells for co-op shares in apartment buildings and row houses. The key information for our analysis includes unique IDs for buyers, sellers, and apartments, debt-free prices, ownership shares, transaction dates, and information on whether the owner received the apartment free of charge (e.g. inheritance).

Table A1: Data sources

Data	Source	Years
Micro-level		
Transfer tax register	Statistics Finland (StatFi)	2006-2018
Apartment register	StatFi	2005-2018
Property register	StatFi	2005-2018
Household information (FOLK-perus)	StatFi	2005-2018
Residence data	StatFi	2005-2018
Housing loan balance	StatFi	2005-2018
Classifications and other		
Municipality central locations	National Land Survey of Finland	2018
Travel-to-work area classification	StatFi	2018
Real housing price index (100=2005)	StatFi	2005-2018
Price Area classification	StatFi	2015

[22] In the Finnish system, individuals and companies must report purchases of real estate, apartments in a housing company, or other securities or corporate stocks.

We complement the transfer tax data with apartment, property, and household characteristics from administrative data provided by Statistics Finland. The characteristics of the apartments and properties are used in the hedonic price models. The apartment characteristics include, for example, room count and floor area. The property characteristics include, for example, the age of the building, number of floors, building type, building location in 250 x 250 grids, and dummies for whether the property has an elevator and whether the housing company owns the plot under the building.

Household characteristics are used as controls in the main analysis. The variables include owner age, disposable income, housing debt, and dummies for whether the household has children, two owners and higher education. Where there are two owners, these characteristics are aggregated at the household level. End-of-year residence information allows us to detect moves and restrict our analysis to homeowners.

A1.1.2 Sample selection

The transfer tax data cover the years 2006 to 2018. The data are truncated below from 2006 meaning that we observe ownership spells that started before 2006 if they were still ongoing in 2006. For these spells, we observe all relevant information on purchases such as the debt-free purchase price, purchase date, and ownership shares. The seller ID, however, is not observed for spells that started before 2006. From the above, the data are censored from 2018 meaning that we can observe the termination of an ownership spell only if it occurred before 2019. From the data we observe all market transactions that occurred between 2006 and 2018.

The raw transfer tax data are organized at the individual owner level. This means that if a couple jointly own an apartment in a given year, they both have separate entries in the data with a specific variable describing the ownership share of that individual. In the analysis, however, we are interested in ownership spells of households rather than individuals. Before converting the data to the apartment level, we excluded those apartments that had missing or multiple values for some of the key variables such as spell level sales dates or prices. We also excluded apartments that had a nontypical transaction history such as less than 6 month ownership spells or company involved, as these apartments have likely experienced unobserved quality improvements (e.g. house flips) that make their market price difficult to model correctly. Additionally, we excluded apartments that were received free of charge (e.g. as bequests), that had more than two owners at some point in time, and those that had been only partly sold. This is to make sure that the realized returns correspond to the homeowners' market returns.

From the apartment-level data, we excluded all the apartments that did not match Statistics Finland's apartment or property registers and apartments that had missing or multiple values in key explanatory variables. Other restrictions included dropping detached houses, apartments with floor areas less than 20 or more than 150 square meters, apartments with more than 6 rooms, and those built or first purchased before 1900. To this data, we applied separate additional sets of restrictions to reach the two data sets required for the analysis. The floor area consistency rule differs slightly between the two: for the matched apartment-household panel we require that the recorded floor area does not vary by more than 3 square meters across years, to accommodate minor measurement revisions; for the transaction data we require a single consistent floor area value across all observations for the apartment.

For the matched apartment-household panel the additional restrictions include, for example, dropping apartments that had a mismatch between sales and purchase transactions, or sales and purchase price. Also, apartments that were in the top 0.5 percent in the distribution of debt-free price or nominal return in any year of the panel in a given price area were dropped. Finally, we excluded ownership spells that did not match Statistics Finland's registers of apartment coordinates or owners' background information, had missing years in ownership spells, had underaged owners, and those where the owner died during the spell. In these data, we classified apartments as owner-occupied in year t if (both) owner(s) lived in the apartment at the end of $t-1$. A move occurs when the owner(s) change(s) address(es).

For the transaction data, the additional restrictions included only dropping transactions that were in the top or bottom 0.5 percent in the distribution of debt-free price in a given price area and year.

A1.1.3 Summary statistics

Table A2: Summary statistics

	Helsinki TWA	Tampere TWA	Turku TWA
Sales	67008	21842	13421
Apartments	108769	33535	23151
Debt-free sales price (2005 EUR)	188155 (85744)	134117 (53422)	123781 (48819)
Floor-area	69 (23)	69 (20)	72 (21)
Room count	3 (1)	3 (1)	3 (1)
Sauna	0.49	0.61	0.49
Balcony	0.69	0.66	0.5
Separate kitchen	0.78	0.85	0.87
Floor count	4 (2)	3 (2)	4 (3)
Year of construction	1977 (19)	1980 (16)	1977 (16)
Own lot	0.76	0.73	0.92
Elevator	0.44	0.38	0.51
Row-house	0.28	0.38	0.34
Distance to CBD	5.62	4.28	2.95

Notes: The table reports summary statistics for apartment characteristics for the full panel before bandwidth restrictions. N = 1,752,331

A1.2 Additional Results

A1.2.1 Price model

Table A3: Test sample predictive performance, 2006-2018

		Share of observations with absolute relative residual			
Area	Method	< 0.05	< 0.10	< 0.15	< 0.20
Helsinki TWA	Ensemble	0.441	0.725	0.866	0.930
Helsinki TWA	OLS	0.310	0.568	0.741	0.850
Helsinki TWA	Bagging	0.424	0.694	0.840	0.915
Helsinki TWA	Random forest	0.425	0.705	0.851	0.921
Helsinki TWA	XGBoost	0.433	0.718	0.863	0.928
Tampere TWA	Ensemble	0.417	0.693	0.838	0.905
Tampere TWA	OLS	0.302	0.564	0.739	0.848
Tampere TWA	Bagging	0.413	0.673	0.818	0.892
Tampere TWA	Random forest	0.408	0.679	0.825	0.896
Tampere TWA	XGBoost	0.406	0.684	0.832	0.904
Turku TWA	Ensemble	0.363	0.645	0.803	0.885
Turku TWA	OLS	0.249	0.473	0.644	0.771
Turku TWA	Bagging	0.350	0.614	0.779	0.867
Turku TWA	Random forest	0.353	0.619	0.783	0.871
Turku TWA	XGBoost	0.361	0.635	0.801	0.886

Notes: The table reports the share of test sample observations with absolute relative residual, defined as $|1 - \exp(\hat{p} - p)|$, below each threshold by model.

Table A4: Test sample performance(RMSE)

Method	Helsinki TWA	Tampere TWA	Turku TWA
Ensemble	**0.124**	**0.139**	**0.144**
OLS	0.159	0.168	0.187
Bagging	0.132	0.148	0.153
Random forest	0.128	0.143	0.149
XGBoost	0.125	0.141	0.145

Notes: RMSE (root mean squared error) is calculated as $\text{RMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(p_i - \hat{p}_i)^2}$, where p_i is the observed log-price and $\hat{p}_i$ is the predicted log-price.

Table A5: Ensemble weights by area

Method	Helsinki TWA	Tampere TWA	Turku TWA
OLS	0	0.012	0
Bagging	0.023	0.056	0.032
Random forest	0.341	0.311	0.351
XGBoost	0.636	0.621	0.617

Ensemble weights $\hat{w}$ are estimated from the training sample sample by $\min_{w\geq 0} \|p - Xw\|^2$, where p is the observed log-price and X contains the out-of-fold model predictions for OLS, Bagging, Random forest and XGBoost.

Table A6: Predictive performance (RMSE), Helsinki TWA

Method	Training sample (CV)	Training sample	Test sample	Ensemble weight
Ensemble		**0.122**	**0.124**	
OLS	0.157	0.155	0.159	0
Bagging	0.129	0.062	0.132	0.023
Random forest	0.125	0.074	0.128	0.341
XGBoost	0.124	0.108	0.125	0.636

Table A7: Predictive performance (RMSE), Tampere TWA

Method	Training sample (CV)	Training sample	Test sample	Ensemble weight
Ensemble		**0.138**	**0.139**	
OLS	0.167	0.164	0.168	0.012
Bagging	0.145	0.071	0.148	0.056
Random forest	0.14	0.085	0.143	0.311
XGBoost	0.14	0.112	0.141	0.621

Table A8: Predictive performance (RMSE), Turku TWA

Method	Training sample (CV)	Training sample	Test sample	Ensemble weight
Ensemble		**0.146**	**0.144**	
OLS	0.19	0.186	0.187	0
Bagging	0.153	0.073	0.153	0.032
Random forest	0.149	0.089	0.149	0.351
XGBoost	0.148	0.113	0.145	0.617

A1.2.2 Predicted returns

Figure A1: Panel observations by expected loss/gain in Helsinki, Tampere and Turku TWAs

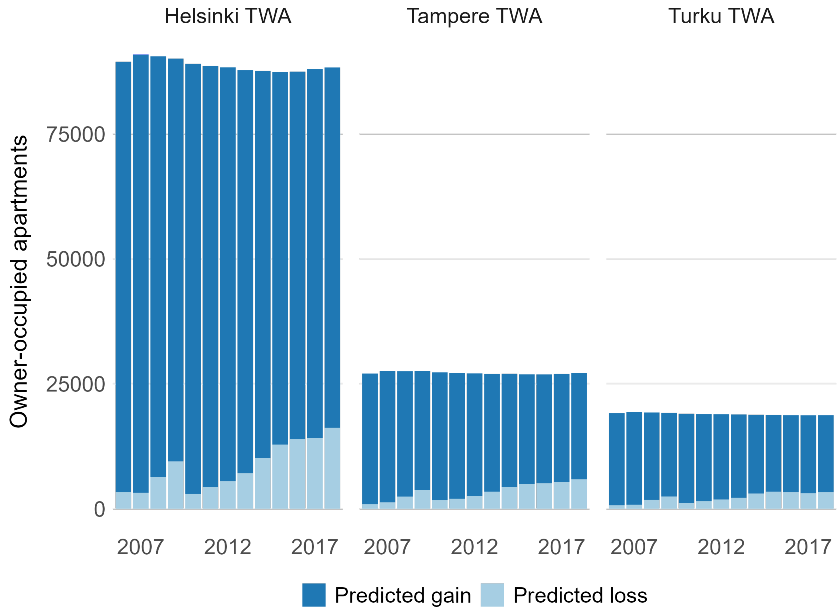

Figure A2: Missing and excess mass in the nominal capital gains distribution

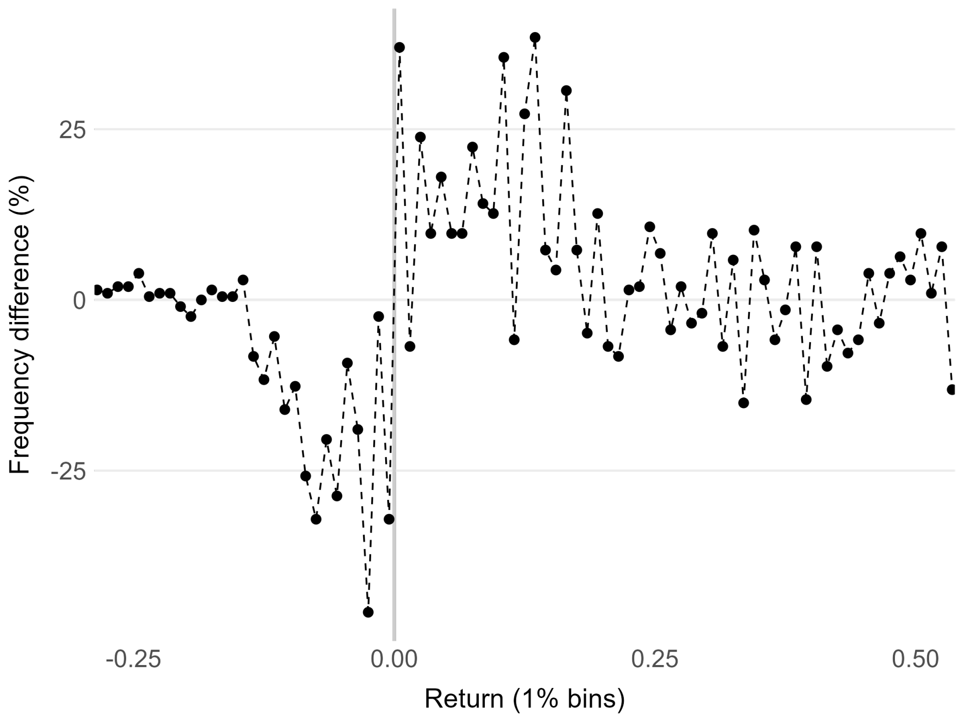

Notes: Figure shows the difference between realized and predicted nominal frequencies of returns, binned in 1% intervals for the realized transaction data. Positive values indicate bins where the realized frequencies exceed the predicted frequencies. Vice versa for negative values.

A1.2.3 Additional results

Figure A3: Mobility and employment

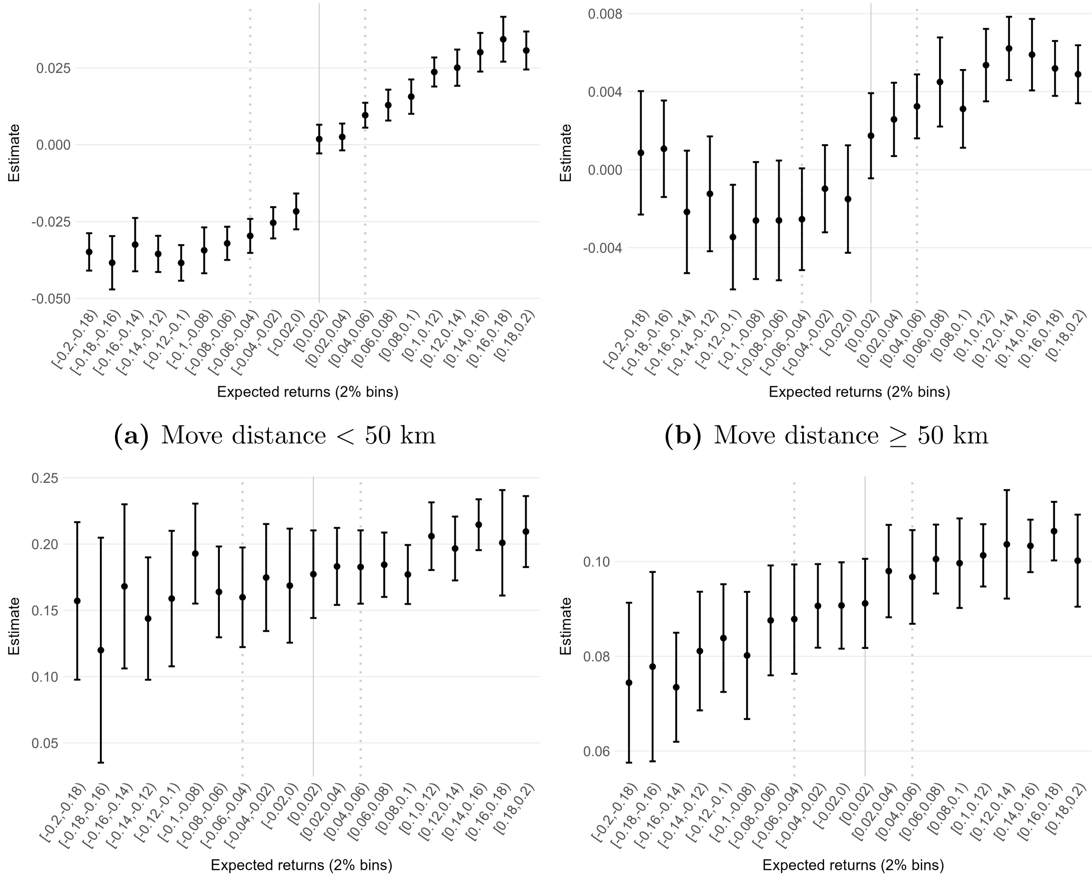

(a) Move distance < 50 km

(b) Move distance ≥ 50 km

(c) End of unemployment

(d) Change of employer

Notes: Figure shows estimates ($\hat{\beta}_i$) for expected return bins (2% intervals) from regressions of the form $Y_{it} = \sum_i \beta_i \, \mathrm{Bin}_i + \gamma \, X_{it} + \delta \, \mathrm{Year} + \varepsilon_{it}$, where X include controls for spell length (entered in linear and squared form) and loan-to-value ratio (in 5% bins). Solid vertical lines indicate the bin containing zero expected return. Dotted vertical lines indicate $\pm$ 5% expected return, which is used as an omitted (donut) region in our main specification. Error bars show 95% confidence intervals based on standard errors clustered at the municipality level.

Table A9: Main regression results - coefficients on expected return

Outcome	Obs	Return	LOSS x Return
Sale	386,698	0.216 *** (0.015)	-0.197 *** (0.025)
Move	386,698	0.235 *** (0.015)	-0.212 *** (0.024)
Sale conditional on move (t+1)	34,076	-0.029 (0.035)	0.674 *** (0.214)
No sale conditional on move (t+1)	34,076	0.029 (0.035)	-0.674 *** (0.214)
Became landlord conditional on move	34,076	0.042 (0.049)	-0.312 (0.192)
Owner-occupied conditional on move	34,076	0.062 (0.058)	0.437 * (0.218)
Rental conditional on move	34,076	0.015 (0.039)	-0.342 * (0.176)
Move within TWA	386,698	0.213 *** (0.016)	-0.165 *** (0.024)
Move between TWAs	386,698	0.019 *** (0.006)	-0.048 *** (0.014)
Move distance less than 50 km	386,698	0.204 *** (0.016)	-0.159 *** (0.024)
Move distance over 50 km	386,698	0.019 *** (0.006)	-0.043 *** (0.012)
End of unemployment	26,185	0.235 *** (0.056)	0.027 (0.154)
Move and end of unemployment	26,185	0.121 ** (0.051)	-0.141 *** (0.042)
Move and end of unemployment, over 50km	26,185	0.038 *** (0.013)	-0.060 *** (0.022)
Change of employer	304,703	0.033 ** (0.014)	0.043 (0.041)
Move and change employer	304,703	0.071 *** (0.008)	-0.055 *** (0.011)
Move and change employer, over 50km	304,703	0.008 (0.005)	-0.014 (0.009)

Notes: The table reports estimates for β and γ from equation (1): $Y_{it} = \alpha_t + \tau LOSS_{it} + \beta Return_{it} + \gamma(LOSS_{it} \times Return_{it}) + \mu X_{it} + \varepsilon_{it}$. The estimation uses a bandwidth of -20% to 20% around the zero return cutoff, excluding observations in the donut hole $[-5\%, 5\%]$. Standard errors clustered at the municipality level are in parentheses. Significance is denoted by asterisks: *p<.1,**p<.05 ,***p<.01.

Table A10: Persistence of landlord status and no-sale

Outcome	Obs	Control mean	Loss
No sale conditional on move (t+1)	26,597	0.128	0.100 *** (0.026)
No sale conditional on move (t+2)	26,597	0.107	0.071 ** (0.029)
No sale conditional on move (t+3)	26,597	0.094	0.051 * (0.027)
Became landlord conditional on move	26,597	0.102	0.061 ** (0.024)
Post-move landlord at t+2	26,597	0.081	0.084 *** (0.021)
Post-move landlord at t+3	26,597	0.073	0.038 * (0.023)

Notes: The table reports estimates for τ from equation (1): $Y_{it} = \alpha_t + \tau LOSS_{it} + \beta Return_{it} + \gamma(LOSS_{it} \times Return_{it}) + \mu X_{it} + \varepsilon_{it}$. $LOSS_{it}$ and $Return_{it}$ use fitted values from a hedonic price model estimated by the ensemble model. The estimation uses a bandwidth of -20% to 20% around the zero return cutoff, excluding observations in the donut hole $[-5\%, 5\%]$. The control mean corresponds to the predicted outcome just above the zero return cutoff with control variables X_{it} set to their sample means. Estimates for β and γ are reported separately in Table A9. Standard errors clustered at the municipality level are in parentheses. Significance is denoted by asterisks: *p<.1,**p<.05 ,***p<.01.

Figure A4: Heterogeneity: Sale

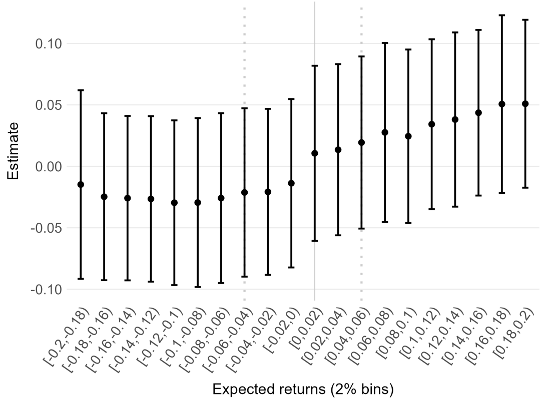

(a) Low LTV, Short spell

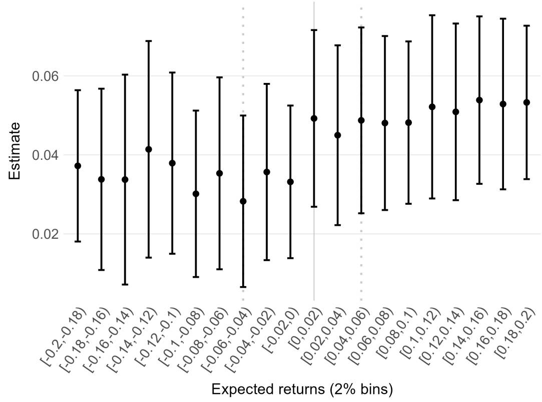

(b) Low LTV, Long spell

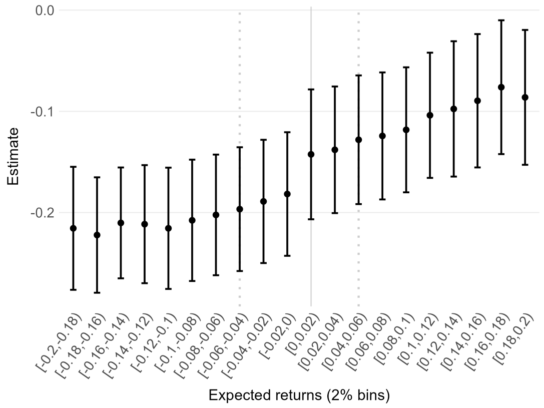

(c) High LTV, Short spell

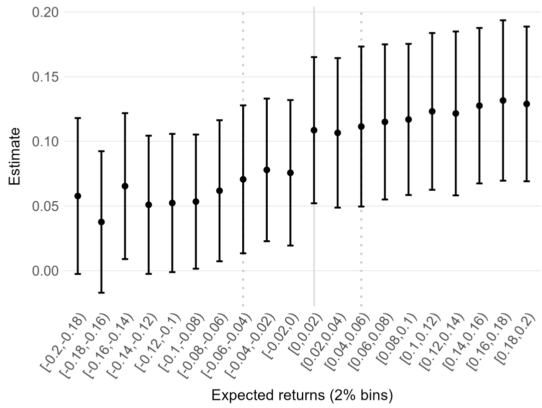

(d) High LTV, Long spell

Notes: Figure shows estimates ($\hat{\beta}_i$) for expected return bins (2% intervals) from regressions of the form $Y_{it} = \sum_i \beta_i \,\mathrm{Bin}_i + \gamma\, X + \delta\, \mathrm{Year} + \varepsilon$, fitted to sub-samples based on LTV and spell length. X include controls for spell length (entered in linear and squared form) and loan-to-value ratio (in 5% bins). Solid vertical lines indicate the bin containing zero expected return. Dotted vertical lines indicate ± 5% expected return, which is used as an omitted (donut) region in our main specification. Error bars show 95% confidence intervals.

Figure A5: Heterogeneity: No-sale conditional on move (t+1)

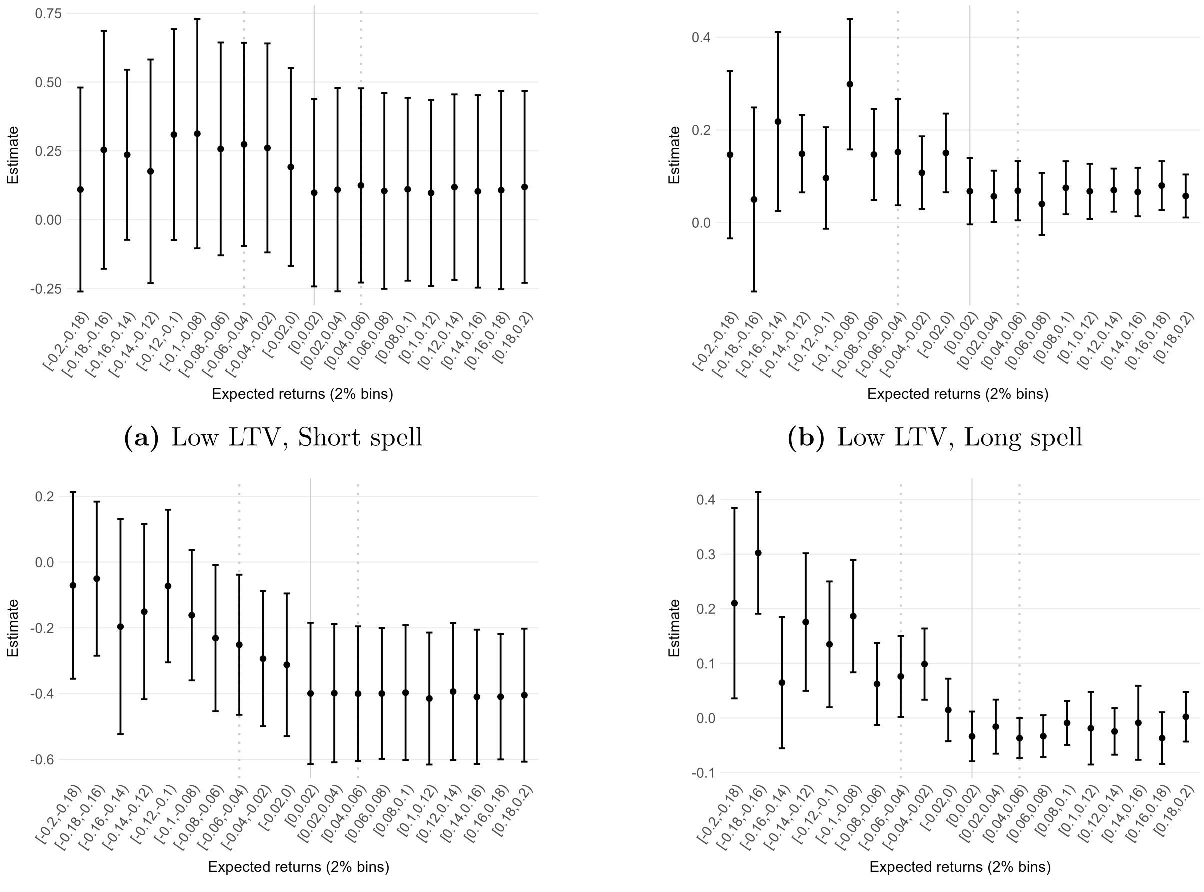

(a) Low LTV, Short spell

(b) Low LTV, Long spell

(c) High LTV, Short spell

(d) High LTV, Long spell

Notes: Figure shows estimates ($\hat{\beta}_i$) for expected return bins (2% intervals) from regressions of the form $Y_{it} = \sum_i \beta_i \, \text{Bin}_i + \gamma \, X + \delta \, \text{Year} + \varepsilon$, fitted to sub-samples based on LTV and spell length. X include controls for spell length (entered in linear and squared form) and loan-to-value ratio (in 5% bins). Solid vertical lines indicate the bin containing zero expected return. Dotted vertical lines indicate ± 5% expected return, which is used as an omitted (donut) region in our main specification. Error bars show 95% confidence intervals.

Figure A6: Heterogeneity: Became landlord on move (t+1)

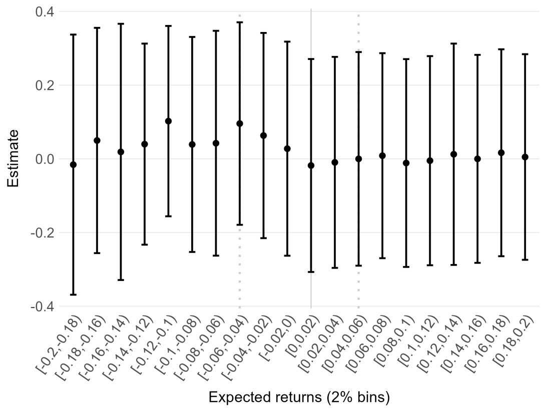

(a) Low LTV, Short spell

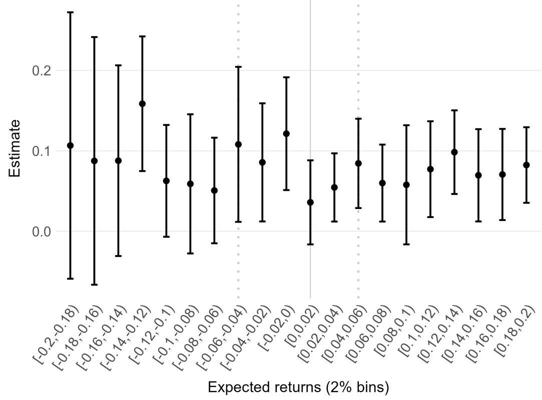

(b) Low LTV, Long spell

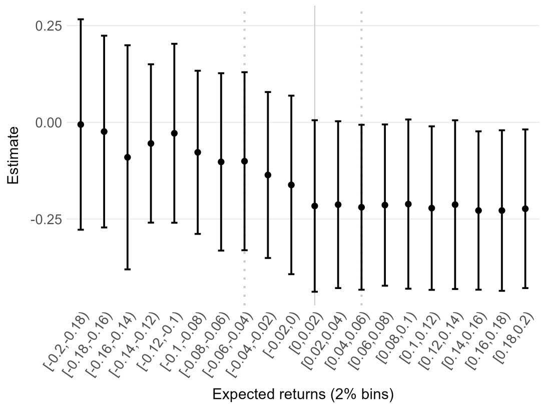

(c) High LTV, Short spell

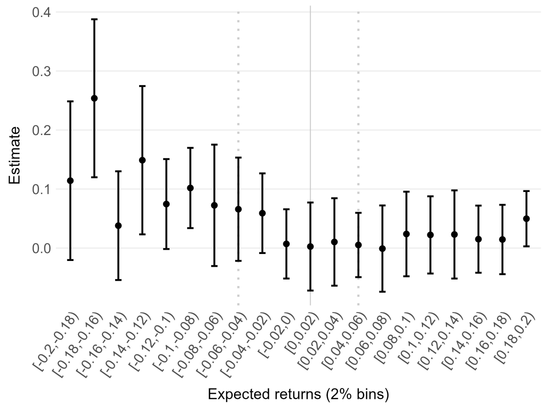

(d) High LTV, Long spell

Notes: Figure shows estimates ($\hat{\beta}_i$) for expected return bins (2% intervals) from regressions of the form $Y_{it} = \sum_i \beta_i \,\mathrm{Bin}_i + \gamma\, X + \delta\, \mathrm{Year} + \varepsilon$, fitted to sub-samples based on LTV and spell length. X include controls for spell length (entered in linear and squared form) and loan-to-value ratio (in 5% bins). Solid vertical lines indicate the bin containing zero expected return. Dotted vertical lines indicate ± 5% expected return, which is used as an omitted (donut) region in our main specification. Error bars show 95% confidence intervals.

Figure A7: Heterogeneity: Rental conditional on move

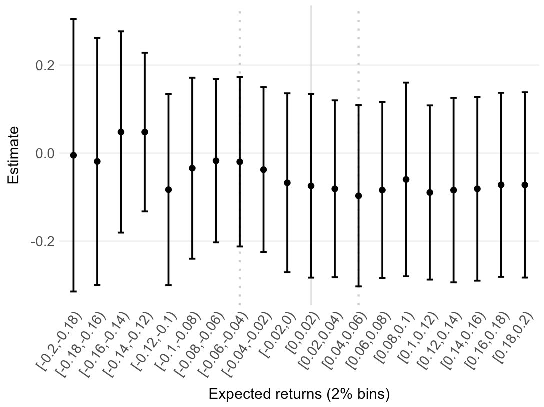

(a) Low LTV, Short spell

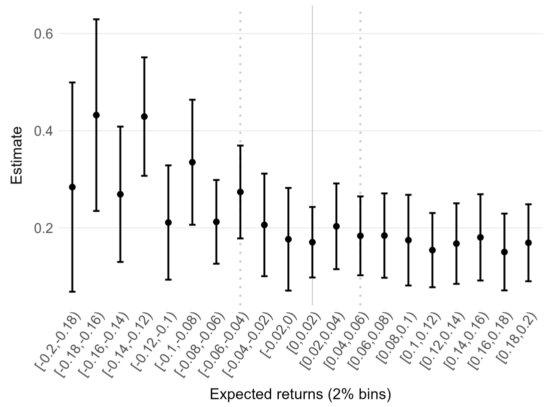

(b) Low LTV, Long spell

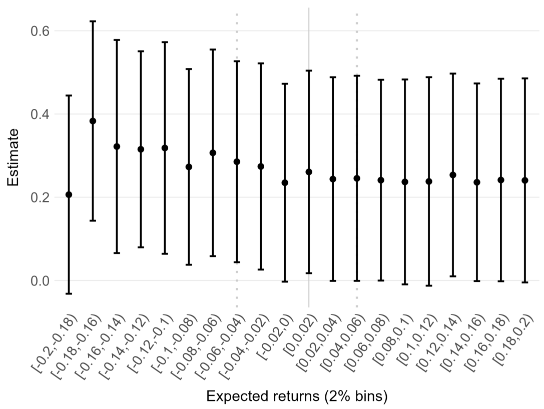

(c) High LTV, Short spell

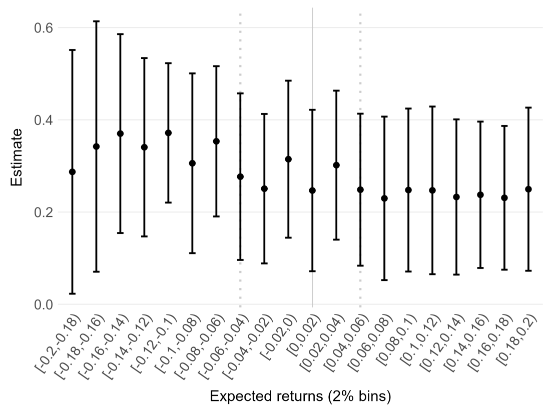

(d) High LTV, Long spell

Notes: Figure shows estimates ($\hat{\beta}_i$) for expected return bins (2% intervals) from regressions of the form $Y_{it} = \sum_i \beta_i \, \text{Bin}_i + \gamma \, X + \delta \, \text{Year} + \varepsilon$, fitted to sub-samples based on LTV and spell length. X include controls for spell length (entered in linear and squared form) and loan-to-value ratio (in 5% bins). Solid vertical lines indicate the bin containing zero expected return. Dotted vertical lines indicate ± 5% expected return, which is used as an omitted (donut) region in our main specification. Error bars show 95% confidence intervals.

A1.2.4 Validity and robustness

Table A11: Robustness checks

Outcome	Ensemble	OLS	With covariates	Spell length	Second-order
Sale	-0.512***	-0.582***	-0.494***	-0.512***	-0.288***
Move	-0.435***	-0.471***	-0.391***	-0.435***	-0.239***
Sale conditional on move (t+1)	-0.137***	-0.111***	-0.134***	-0.136***	-0.041
No sale conditional on move (t+1)	0.714***	0.66***	0.755***	0.708***	0.211
Became landlord conditional on move	0.612***	0.762***	0.67***	0.609***	0.672
Owner-occupied conditional on move	-0.055**	-0.064***	-0.047	-0.055**	-0.008
Rental conditional on move	0.341**	0.286***	0.312**	0.341**	-0.071
Move within TWA	-0.395***	-0.459***	-0.342***	-0.382***	-0.237***
Move between TWAs	-0.667***	-0.462***	-0.583***	-0.667***	-0.167
Move distance less than 50 km	-0.392***	-0.463***	-0.351***	-0.392***	-0.216***
Move distance over 50 km	-0.727***	-0.5***	-0.636***	-0.727***	-0.182
End of unemployment	0.056	0.123***	0.077	0.058	0.034
Move and end of unemployment	-0.467**	-0.396***	-0.452**	-0.467**	-0.333
Move and end of unemployment, over 50km	0.167	0.429*	0.167	0.167	2.167
Change of employer	-0.047	-0.018	-0.039	-0.047	-0.058
Move and change employer	-0.455***	-0.5***	-0.474***	-0.455***	-0.143
Move and change employer, over 50km	-0.6***	-0.6***	-0.75***	-0.6***	-0.4

Notes: The table reports estimates for τ relative to control mean from equation (1) with two alternative price estimates (Ensemble and OLS) and with specifications using alternative sets of controls. The estimation uses a bandwidth of -20% to 20% around the zero return cutoff, excluding observations in the donut hole $[-5\%, 5\%]$. The control mean corresponds to the predicted outcome just above the zero return cutoff with control variables X_{it} set to their sample means. Significance is denoted by asterisks: *p<.1, **p<.05 , ***p<.01.

Figure A8: Balance of household characteristics

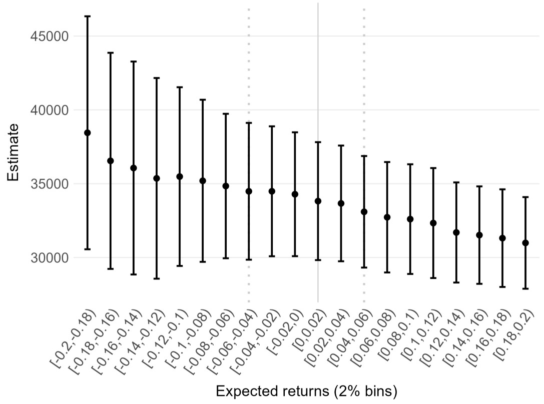

(a) Disposable income

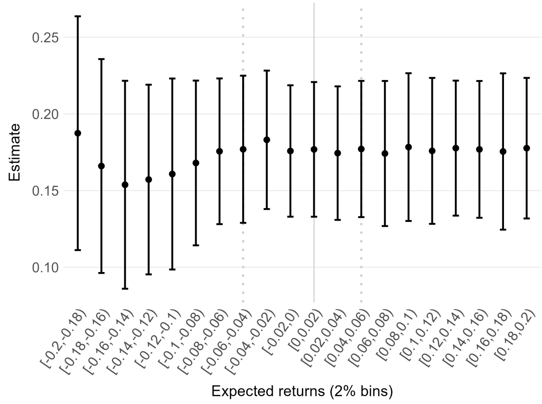

(b) Higher education

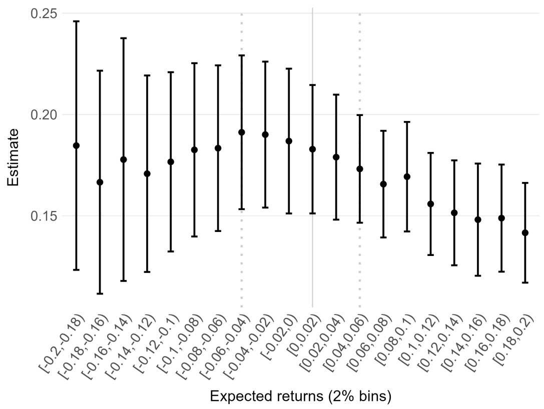

(c) Kids

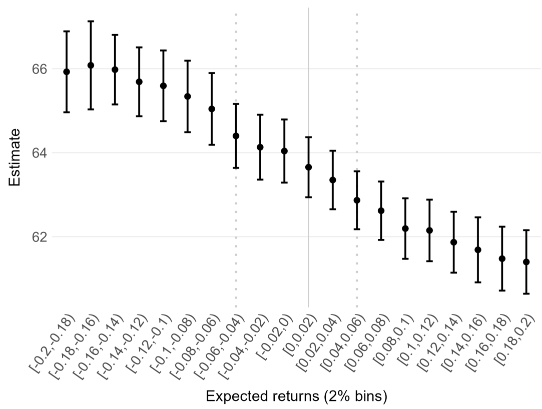

(d) Mean owner age

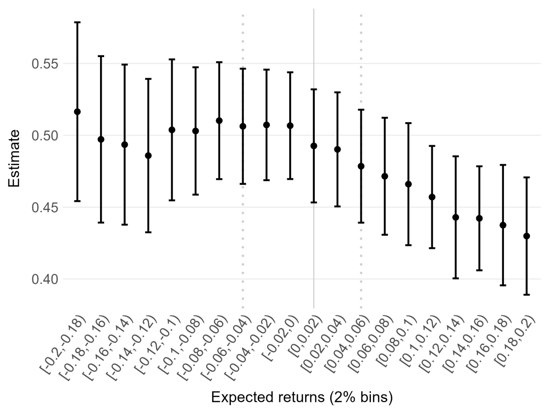

(e) Two owners

Notes: Figure shows estimates ($\hat{\beta}_i$) for expected return bins (2% intervals) from regressions of the form $Y_{it} = \sum_i \beta_i \, \mathrm{Bin}_i + \gamma X + \delta \, \mathrm{Year} + \varepsilon$, where X include controls for spell length (entered in linear and squared form) and loan-to-value ratio (in 5% bins). Solid vertical lines indicate the bin containing zero expected return. Error bars show 95% confidence intervals.

Figure A9: Tenure type of old and new apartment

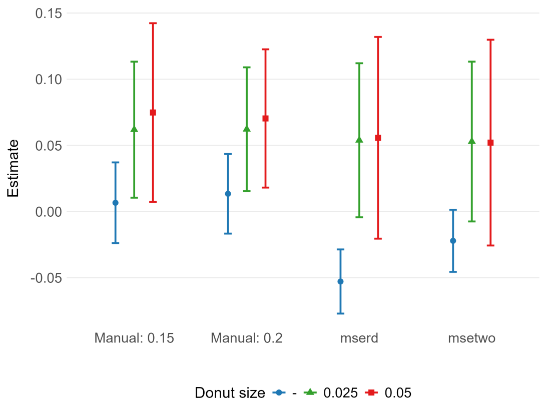

(a) Became landlord conditional on move

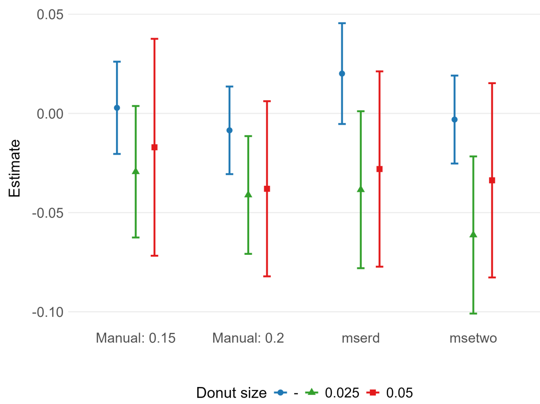

(b) Owner-occupier conditional on move

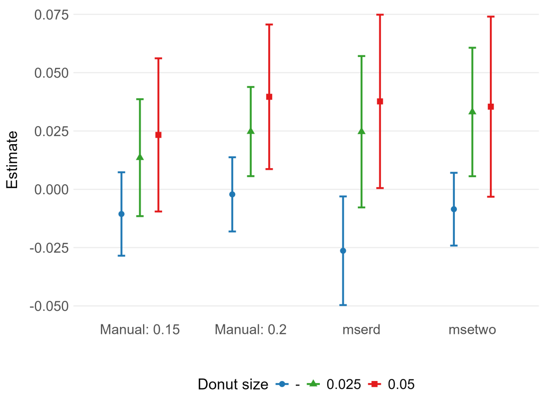

(c) Rental conditional on move

Notes: The figure shows local polynomial estimates of gains with conventional CIs for manual bandwidth models and bias-corrected estimates of gains with robust CIs for MSE-optimal models (Calonico et al., 2014). "mserd" uses symmetric bandwidths, while "msetwo" allows asymmetric bandwidths. All models use first-order polynomials, uniform kernel weighting, and control for LTV dummies, spell length, and year dummies.

Figure A10: Distance of the move

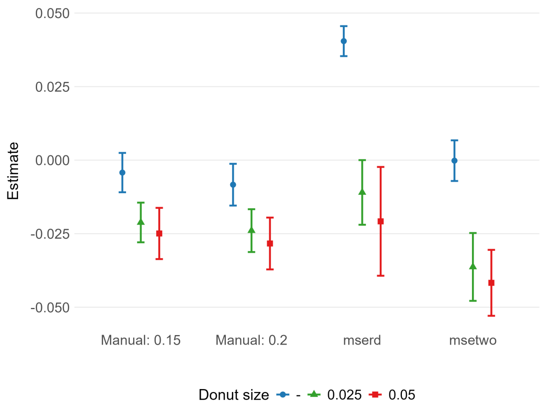

(a) Move distance < 50 km

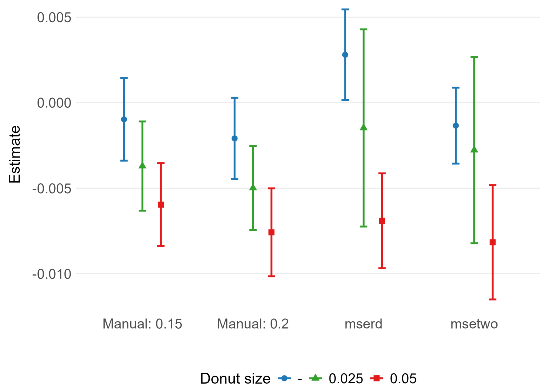

(b) Move distance ≥ 50 km

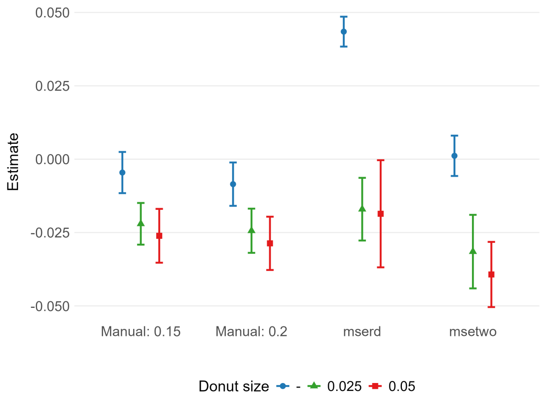

(c) Move within TWA

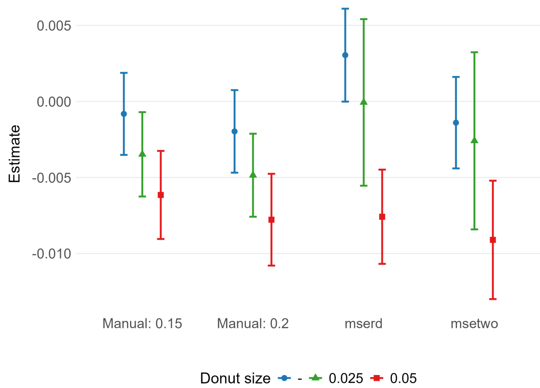

(d) Move between TWAs

Notes: The figure shows local polynomial estimates of gains with conventional CIs for manual bandwidth models and bias-corrected estimates of gains with robust CIs for MSE-optimal models (Calonico et al., 2014). "mserd" uses symmetric bandwidths, while "msetwo" allows asymmetric bandwidths. All models use first-order polynomials, uniform kernel weighting, and control for LTV dummies, spell length, and year dummies.

Figure A11: Employment

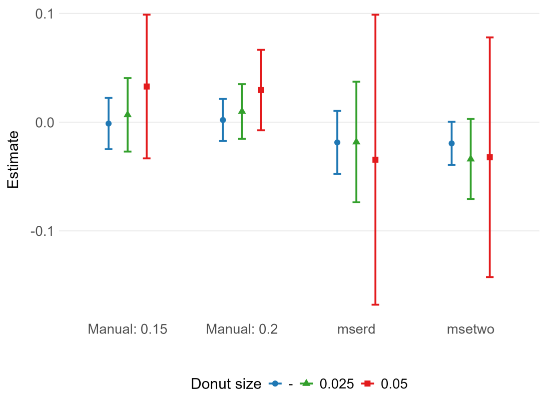

(a) End of unemployment

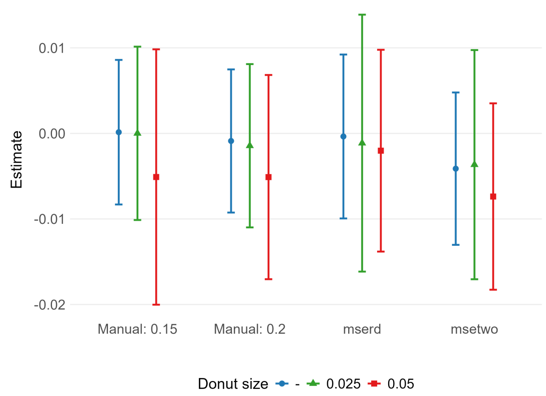

(b) Change of employer

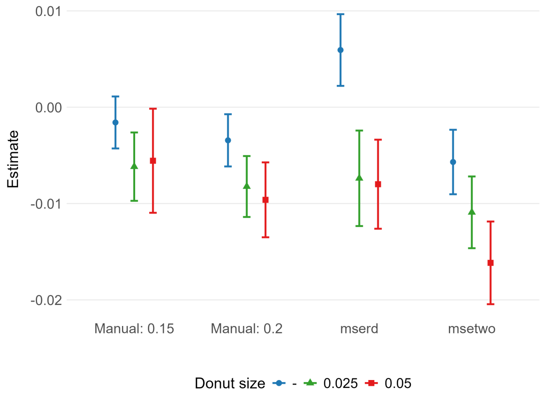

(c) Move and change employer

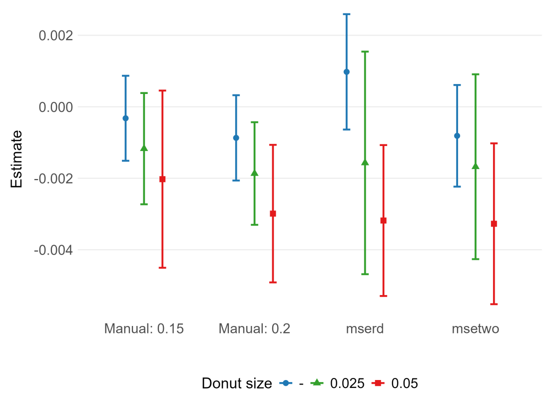

(d) Move and change of employer, over 50km

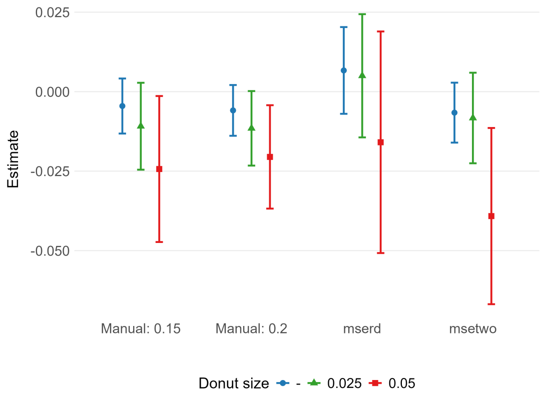

(e) Move and end of unemployment

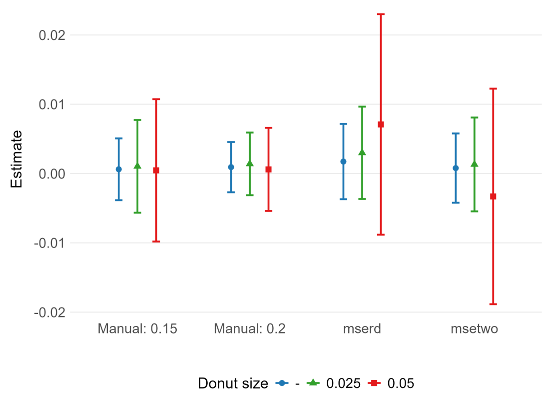

(f) Move and end of unemployment, over 50km

Notes: The figure shows local polynomial estimates of gains with conventional CIs for manual bandwidth models and bias-corrected estimates of gains with robust CIs for MSE-optimal models (Calonico et al., 2014). "mserd" uses symmetric bandwidths, while "msetwo" allows asymmetric bandwidths. All models use first-order polynomials, uniform kernel weighting, and control for LTV dummies, spell length, and year dummies.