Seller and Buyer Behavior in the Housing Market

Selger- og kjøperadferd i boligmarkeder

Philosophiae Doctor (PhD) Thesis

Andreas Eidspjeld Eriksen

Norwegian University of Life Sciences
School of Economics and Business

Ås (2024)

Thesis number 2025:7
ISSN 1894-6402
ISBN 978-82-575-2216-2

Supervisors and Evaluation Committee

Supervisor:

Dag Einar Sommervoll, Norwegian University of Life Sciences

Co-supervisors:

André Kallåk Anundsen, Oslo Metropolitan University

Cloé Garnache, Oslo Metropolitan University and University of Oslo

Plamen Nenov, Norges Bank

Evaluation Committee:

Acknowledgements

It has been a true privilege to have the opportunity to write this thesis, which has allowed me to cultivate my research interests. While at times I have found myself being stressed and in doubt, this four-year journey has been a pleasure, both during the process and in hindsight.

Most importantly, I want to thank my supervisors André Kallåk Anundsen, Cloé Garnache, Plamen Nenov, and Dag Einar Sommervoll for their guidance, support, and knowledgeable advice. Their optimism, dedication, ambition, and genuine interest have helped me stay motivated and maintain joy in my efforts to complete the thesis. Cloé has been unwavering throughout the process, and I have enjoyed working with her and getting to know her. I truly appreciate her creativity, experience, and infectious dedication to research. André was mostly involved during the first half of my PhD period, providing clear and valuable feedback to my papers and presentations, which have greatly helped me later on. I am grateful to Plamen for accepting to be my co-supervisor when André left academia, particularly for his help in directing my focus in my paper about buyer search. I also want to thank Dag Einar for his support, helpful comments, and general feedback.

I have been fortunate to work with great colleagues. Although we have been a small group, the work environment at Housing Lab has been stimulating and genuinely positive, much thanks to the efforts of Erling Røed Larsen. I extend my gratitude to him for believing in me when I first reached out hoping he would supervise me in writing my master thesis. Throughout my years as a PhD student, Erling has always found time to help, listen, provide valuable feedback and suggestions, and offer general support, even though he has not been my supervisor. I want to thank my fellow PhD students Nini Barth and Jeanette Fjære-Lindkjenn for struggling together with me in the early phase of the PhD process. I also want to thank Andreas Benedictow, Bjørnar Karlsen Kivedal, and Erlend Eide Bø for fruitful discussions, constructive feedback, and contributions.

In writing this thesis, I have spent may hours outside of the usual working schedule, including long days, evenings, weekends, and holidays. I extend my sincere gratitude to my wife, Karoline, for her sustained support and understanding throughout these years. I want to thank my son, Alfred, who was born late in the second year of writing, for understanding that his dad often needed to work. In addition, I want to thank my soon-to-be-born twin daughters, who have motivated me to finish this thesis. Finally, I want to thank my family for their understanding and support.

Contents

List of papers

Paper 1: Eriksen, A. E. (2024): On Loss Aversion in Housing Markets: Evidence from Norway.

Paper 2: Eriksen, A. E. and Garnache, C. (2024): Here Comes the Sun: The Effect of Sunshine on Home Prices.

Paper 3: Eriksen, A. E. (2024): How directed is housing search?

Paper 4: Anundsen, A. K., Benedictow, A., Eriksen, A. E., Røed Larsen, E., and Walbækken, M. (2024): Repeat bidder behavior in housing auctions.

Abstract

This thesis examines the behavior of sellers and buyers in the housing market across different scenarios and perspectives. Individuals engaging in the housing market are often inexperienced, therefore, they may potentially be prone to acting in accordance with heuristics and cognitive biases. This in turn can affect market outcomes. Because this market is a large part of the economy, such micro-level behavior impacts the macroeconomy. The overall findings suggest that individuals are affected by these influences.

The first paper investigates whether sellers in the housing market act in accordance with loss aversion. The paper estimates the effect of loss aversion on sellers' list price choices by following an established identification method and proposes a new approach to address the known identification problem. This problem relates to measuring price expectations. The results suggest that the identification method is more sensitive to omitted variables bias than previously believed, and that sellers' exhibition of loss aversion likely depends on institutional factors.

The second paper (with Garnache) investigates whether sunlight affects housing prices. Sunlight can affect how buyers form beliefs, impressions, and decisions about housing units of interest, potentially through misattribution of mood. There are two main events for sunlight to directly affect the buyers: the public showings and the date of the agreement to sell. The results suggest that sunlight at the showings affect price formation, but not at the day of the sale. The sunlight effect is more pronounced when the days are shorter and darker.

The third paper explores whether list prices are important determinants of buyer search in the housing market. The paper is motivated by the search theory literature, which distinguish between models that assume random and directed search. From a directed search perspective, list prices should be important beyond their capacity to nest information about the housing units for sale. The results suggest that list prices are in general of lesser importance, except for the lower-priced housing segment, where list prices impact buyer search in accordance with directed search.

The fourth paper (with Anundsen, Benedictow, Røed Larsen, and Walbækken) examines how bidders change their bidding behavior across housing unit auctions. The results suggest that bidders change their behavior over time. Specifically, bidders who win their last auction adjust their behavior in line with increasing their likelihood of winning, more so than those consistently losing. Implying, given a fixed budget, those winning are adjusting to increase their competitiveness at a faster rate than those losing.

Norsk sammendrag

Denne avhandlingen undersøker selger- og kjøperadferd i boligmarkedet på tvers av ulike scenarier og perspektiver. Enkeltpersoner som deltar i boligmarkedet er ofte uerfarne i markedet, og de kan derfor være mer tilbøyelige til å handle i tråd med heuristikk og kognitive skjevheter. Dette kan igjen påvirke ulike utfall i markedet. Fordi dette markedet utgjør en stor del av økonomien så påvirker slik mikroadferd makroøkonomien. De samlede funnene antyder at enkeltpersoner påvirkes av slike faktorer.

Den første artikkelen undersøker om selgere i boligmarkedet oppfører seg i tråd med tapsaversjon. Artikkelen estimerer effekten av tapsaversjon på selgeres valg av prisantydning ved å bruke en etablert identifikasjonsmetode, og den presenterer en ny tilnærming for å adressere det kjente identifikasjonsproblemet med denne metoden. Dette problemet er relatert til hvordan prisforventninger måles. Resultatene antyder at identifikasjonsmetoden er mer sensitiv for utelatte variabler enn tidligere antatt, og at selgeres utvisning av tapsaversjon sannsynligvis avhenger av institusjonelle faktorer.

Den andre artikkelen (med Garnache) undersøker om sollys på visninger og salgsdager påvirker boligpriser. Kjøperes oppfattelser og beslutninger kan påvirkes av sollys fordi dette kan påvirke deres humør, og humøret kan feiltolkes til å ha meningsfullt innhold for boligkjøp. Resultatene antyder at sollyset på visningene påvirker prisene, men ikke på salgsdagen. Denne effekten er ikke til stede om sommeren når dagene generelt er lysere. Videre viser resultatene at kjøpere av leiligheter, kjøpere som kjøper alene, mindre velstående kjøpere, og førstegangskjøpere er mer påvirket av sollys enn andre.

Den tredje artikkelen undersøker om prisantydninger er viktige for kjøperes søk i boligmarkedet. Artikkelen er motivert av litteraturen om søketeori som skiller mellom modeller som antar tilfeldig og rettet søk. Fra perspektivet til rettet søk så skal prisantydninger ha en rolle utover deres evne til å inneholde informasjon om boligene til salgs på markedet. Resultatene antyder at prisantydninger generelt er mindre viktige for kjøperes søk, med unntak av i det lavere prissegmentet hvor prisantydningen påvirker søk i tråd med rettet søk.

Den fjerde artikkelen (med Anundsen, Benedictow, Røed Larsen, and Walbækken) undersøker hvordan budgivere endrer budgivingsadferd på tvers av boligauksjoner. Resultatene antyder at budgivere endrer adferd over tid: budgivere som vinner den siste auksjonen de deltar i justerer adferd i tråd med å øke sannsynligheten for å vinne, og i større grad enn de som konsekvent taper. Dette impliserer at for et gitt budsjett så vil de som vinner justere seg for å øke sin konkurranseevne raskere enn de som taper.

1 Introduction to the thesis

1.1 Introduction

The residential segment of the real estate market, called the housing market, is a key economic factor, especially in economies where this market is liberalized. The housing market consists of many sub-markets, which may be geographically separated or differentiated by type of housing, such as local markets for apartments. Knowledge about the housing market is important for many reasons, three of which are highlighted in what follows. First, the primary purpose of housing is to provide the homeowner or the tenant with a roof over the head and safety, meaning they serve as homes. Put differently, housing provides utility to the household that occupies the housing unit.

Second, even though housing is crucial for most people, due to market liberalization of the housing market and other related markets such as the credit market, housing is considered an asset. Homeowners and landlords are investors in the market for housing, and due to the high prices of single units, this means that housing often is the largest asset in most household portfolios (Campbell & Cocco, 2007). It is commonly known that house prices are volatile, thus, there exists a real risk in being invested in this market. The lack of diversification in household portfolios makes them highly exposed to this risk. Hence, there is a duality in the importance of housing especially for households that are homeowners. Moreover, because these households are heavily invested into the housing market, their budgets, debt, and savings are closely related to it. This has implications for the macroeconomy, which can be illustrated with Norway as an example. In 2023, the average price of a housing unit was 4.8 million Norwegian kroner, or roughly 480,000 USD, with a total of 88,000 transactions (Statistics Norway, 2024d). These transactions alone amount to a total of 42.2 billion USD. Assuming that this average transaction price is representative for houses in Norway, this implies that the housing stock is worth 1,300 billion USD, which is about 2.5 times the 2023 Norwegian gross domestic product (Statistics Norway, 2024b, 2024c). With the fact that most households own their home, the general population is clearly heavily invested in this market (Statistics Norway, 2018).[1] Hence, what happens in the housing market has implications for macroeconomic factors that involve consumption, debt, and savings. To this extent, the state of the housing market is argued to be a good indicator for future recessions (Leamer, 2007).

Third, not only are households typically invested in the housing market, but most homeowners also keep their homes and are able to change homes because they receive income from their jobs. By allowing for reallocation of housing through a market, this

[1] The housing stock in 2024 was 2.72 million and the population as of January 1, 2024 was 5.55 million (Statistics Norway, 2024a, 2024b).

makes sure that households can relocate to houses and locations that fit them better. For instance, when changing jobs this may require relocating to a new home due to commuting distance, and the choice of changing jobs may itself dependent on the ease of relocating. This continues to be the case even in recent years when remote work has become more common, because many jobs are not suitable for being performed remotely, while other jobs require workers to be physically present during specific days of the week. Although people can sell and buy housing units on the market, housing is considered an illiquid asset (Han & Strange, 2015). Difficulties with relocating, typically related to decline in house prices leading to lower house equity, can make people reconsider their job opportunities and even make them consider jobs at lower levels (Brown & Matsa, 2020). Another side of the story is when there are issues with the local housing market but a relatively high local demand for labor. A well-functioning housing market facilitates for workers' ability to move closer to areas with high labor demand, but when there is a lack of housing supply in these areas the become unavailable to many workers. The lack of efficient allocation of labor through housing has limited productivity growth, and thus the gross domestic product, in the US (Hsieh & Moretti, 2019). In total, there is little doubt that the housing market is important for the labor market.

All papers in this thesis are focused on the housing market. The volatility in house prices can have major impacts on people's budgets and overall finances, thus, it is important to understand how agents in the housing market act when facing different scenarios. To this end, the papers in this thesis are about seller and buyer behavior in the housing market, all with an empirical approach. The first paper follows an established literature about the tendency of sellers acting in accordance with being loss averse in the housing market. The simplest way to explain loss aversion is by the famous definition from the paper in which the term was introduced, namely that "losses loom larger than gains" (Kahneman & Tversky, 1979, p. 279). If sellers act in accordance with loss aversion, this can explain the positive correlation between prices and volume in the housing market, especially during market busts. Contrary to previous studies, the findings suggest that there is no such effect in the analyzed market, namely the Oslo housing market. To explain the result, the paper provides empirical evidence of the importance of different types of unobserved heterogeneity in housing units as a source of estimation bias.

The second paper asks whether house price formation is affected by sunlight. There is an established literature about the effect of weather, including sunlight, as measured by cloud cover, in behavioral finance. The paper contributes to this literature by providing evidence from a much higher priced asset, namely housing. Sunlight, as measured by cloud cover, can affect buyers when they visit the housing units at public showings and when they submit bids to the sellers. The evidence suggests that buyers are affected by sunlight at the showings, and the effect is prominent only when the days are generally shorter and darker.

The last two papers are related to buyer search, and they both utilize a rare dataset that includes information about bidders and their bids, and buyer search outcomes in the form of buyer arrival. A crucial facet of these data is that they comprise sales in the Norwegian housing market. In this market, sales of all types of on-market housing units, most notably non-distressed sales, are finalized by a bidding process closely resembling an open ascending bid auction. This is relevant for both of these papers.

The third paper investigates whether buyer search in the housing market is directed by list prices. Most of the literature consists of theoretical search models, and some of these also provide empirical evidence of search, but they are generally limited by data availability. The contribution to the literature includes new empirical evidence about buyer search by focusing on buyer arrival at listed housing units. The paper investigates the importance of list prices for buyer search by treating list prices as predictors when employing machine learning methods to model buyer arrival. The results suggest that list prices are relatively unimportant for buyer arrival, except in the market segment for cheaper housing units.

The fourth and final paper documents how bidders behave over time by tracing them across housing auctions. The current literature is very sparse on how bidders change their behavior over time when they are unsuccessful in their offers to housing sellers. The evidence shows that bidders adjust to bid on smaller and lower priced housing units, and their bids become more competitive over time when losing auctions. Those observed to win are found to be adjusting to become more competitive at a faster rate than those who consistently lose.

1.2 Literature and concepts

In this section I describe the literature and concepts relevant for the four papers in the thesis.

1.2.1 Loss aversion in the housing market

For decades, topics of behavioral economics have been studied to great extent, but there is still a lot to uncover. Loss aversion is a concept of prospect theory, first introduced by Kahneman and Tversky (1979) and later revised by Tversky and Kahneman (1992), and as the name suggests, it is all about the importance of prospects. There are three elements of prospect theory directly relevant for the first paper in the thesis, which are lucidly summarized by Barberis (2013). First, value is derived directly from gains and losses.

Second, gains and losses are evaluated from a starting point, that being the reference point. The definition of the reference point has been widely discussed in the literature due to the vague formulation in Kahneman and Tversky (1979). Yet, one way to think about the reference point is provided by Kőszegi and Rabin (2006), in which they assume that the reference point is derived from expectations. Although specifically stated as an extreme assumption, which relates to the implications of this assumption, including the potential for influencing the reference point, this still serves as an important step in conceptualizing the reference point. Third is loss aversion, which is that people feel more strongly about losses than gains. Put differently, the absolute value of a loss is larger than the absolute value of an equally sized gain (Kahneman & Tversky, 1979, 1984).

Prior to the formalization provided by Kőszegi and Rabin (2006), Genesove and Mayer (2001) investigated sellers' loss aversion in the Boston condominium market. They introduced a framework to identify loss aversion by measuring sellers' prospective losses as the difference between the purchase price and their expected selling price when they list their housing units on the market. This difference is truncated at zero. Meaning, they impose the assumption that purchase prices paid by the sellers are their reference points, and that their expectations about the market value determine whether they face prospective gains or losses. To identifying the effect of loss aversion on prices, particularly in the list prices, their main problem is measuring these expectations. While it is commonly known that classical measurement error in explanatory variables attenuates coefficient estimates, they show how their framework gives rise to upward bias in their estimates. Genesove and Mayer (2001) are able to partly deal with the bias by estimating a potentially lower bound of the effect, achieved by downward biasing the estimate, and in total, they find a significant effect of loss aversion on list prices while the lower bound estimate in the selling price model is not significant.

The seminal paper of Genesove and Mayer (2001) gave rise to a literature focused on loss aversion in the housing market, such as Engelhardt (2003) who finds that loss aversion is important for mobility.[2] An important extension is provided by Bokhari and Geltner (2011) who include a variable for the prospective gains of the sellers. They apply their extended model to investigate loss aversion in the commercial real estate market. A key point they make is that, in identifying loss aversion, it is important to consider the difference between the coefficient estimates on prospective gains and losses. Implying, prospective gains may have been an omitted variable in the Genesove and Mayer (2001) study. Yet, Bokhari and Geltner (2011) find an effect with a similar magnitude even in a market with professional and experienced sellers.

[2]Other examples are Andersen et al. (2022), Anenberg (2011), Einiö et al. (2008), Lamorgese and Pellegrino (2022), and Zhou et al. (2022). The method also gave rise to studies of behavioral biases in other markets, such as the art market (Beggs & Graddy, 2009).

1.2.2 Weather, sunlight, and prices

While most will argue that temporary variations in weather should not affect asset pricing, there is a large literature in finance that finds evidence to the contrary. There are two strands in this literature of importance for the second paper. First, and most important, are the studies focused on the effect of sunlight on asset prices. As variations in sunlight are highly spatiotemporally dependent, and due to a lack of good measurements, this strand uses cloud cover to capture sunlight. Saunders (1993) was the first to study this, finding that sunlight measured close to Wall Street had a positive relationship with stock exchange returns. Hirshleifer and Shumway (2003) extend on the Saunders study by investigating stock returns in 26 countries, and by providing a potential channel for which sunlight may affect prices: misattribution of mood. Based on their findings, which are consistent with Saunders (1993), they argue that if transaction costs were low, there exists an optimal sunlight-based trade strategy.

Goetzmann and Zhu (2005) goes even more in depth of the phenomenon. They use a panel of individual investors to investigate the relationship between buying and selling stocks and cloud cover (sunlight) at the investor locations. These investors are located in five large US cities and are not considered market makers. The results suggested no effect of cloud cover on investor behavior, yet, they were able to replicate the positive effect of sunlight on stock returns. When controlling for the change in the daily bid-ask spread the sunlight effect becomes very small and not significant. This control variable is intended to control for liquidity in the market, which is related to market makers rather than individual investors. Thus, they conclude that the effect is due to market maker behavior. This study is followed by Goetzmann et al. (2015) who use survey data of institutional investors[3] and find that cloudy weather is related to these being more likely to perceive stocks and indices as overpriced.

The second strand in this literature is about the effect of seasonal depression, and specifically seasonal affective disorder (herafter SAD), on asset prices in financial markets. In short, SAD is a depressive disorder in which the symptoms are seasonally varying, with a peak in symptoms being in the winter (Magnusson, 2000). A common hypothesis is that it is related to people being exposed to less sunlight during the day, implying that the prevalence of SAD should be dependent on latitude. However, the literature review by Mersch et al. (1999) finds that other factors may be more important. Still, SAD is treated by light therapy (Magnusson & Boivin, 2003).

The first study about the effect of seasonal depression on prices was Kamstra et al. (2003) who measure the influence of SAD by a proxy for the hours of night which only takes non-zero values in the fall and winter. To link SAD with asset prices, they argue

[3] Institutional investors are considered being among the market makers.

that depression is related to a higher degree of risk aversion. They find an overall tendency of SAD affecting stock index returns at different locations across the world. This study is complimented by subsequent studies about seasonally varying risk attitudes, such as in treasury market returns (Kamstra et al., 2015) and mutual funds (Kamstra et al., 2017), while Garrett et al. (2005) show how the SAD effect is fully captured by seasonally varying risk premium when modeling asset markets using a conditional CAPM, meaning, the effect cannot be exploited for profit. However, there are critics of Kamstra et al. (2003), especially Jacobsen and Marquering (2008) and Kelly and Meschke (2010). These argue that even when typical seasonal controls are included, the SAD is simply picking up seasonality.

Moving to the housing market literature, Gourley (2021) investigates the effect of temperature and precipitation on house prices in Jefferson County, Colorado. Both of these weather variables may affect prices differently depending on the season. During the winter, precipitation comes as snow when temperatures are below the freezing point, and this might be perceived differently than rain in the summer season.[4] Therefore, Gourley (2021) splits the year crudely into two seasons, winter and summer, and analyzes the weather effects separately in these seasons. Moreover, the analyzed data does not contain information about the different stages that are important for price formation, that being the public showings and the date of the agreement to sell. Instead, the closing date is used, which may be anywhere between a few days to two months after the date of the agreement to sell. Rather, Gourley (2021) uses the weather measured in the month prior to the month of sale to estimate the effect. The findings suggest that during the summer precipitation is important, but not in the winter. In the summer, the direction of the precipitation effect is dependent on the temperature: when temperatures are low precipitation is associated with lower prices, while precipitation during high temperatures is associated with higher prices. When it comes to the effect of temperature, this is only significant in the winter, where lower temperatures lead to higher prices.

In the literature, the study of Gourley (2021) is the closest to the second paper in the thesis, which is concerned about temporary cloud cover realizations on prices. He points at some possible mechanisms that may be important for the effect: opportunity costs and search costs may depend on the weather; curb appeal, meaning how appealing the housing unit looks, which depends on the weather; and the bargaining process between inexperienced and relatively uninformed agents, who then may be prone to be affected by weather. To provide an explanation for the negative temperature effect in the winter Gourley (2021) points at buyers' search costs.

[4]Gourley (2021) labels precipitation in the winter as snow and in the summer as rain. Winter is defined as November to April and summer as May to October.

1.2.3 Sunlight, mood, and prices

As described above, the behavioral finance literature provides an alternative explanation for the effects found by Gourley (2021), namely that buyers may be affected by mood induced by weather, and that the mood ultimately affects their behavior. The second paper in the thesis goes beyond Gourley by estimating the effect of specifically sunlight at the two points in time that is important for the buyers to form beliefs, impressions, and make decisions about the house they are interested in buying: the showings and the date of the agreement to sell. The distinction between the two events is important because the final selling price is decided at the date of the sale, and not at the showings. If sunlight at the showings is important for prices, then what is encoded in the buyers' memories at that time and later recalled at the date of the sale affects their decision making. Meaning, the channel of which sunlight at the showings is affecting prices must be through the recall of memories and the content of those memories.

When considering the role of mood in the relationship between sunlight and prices, the psychology literature about mood and judgment can help uncover relevant mechanisms. Whether it is a sunny day has been documented by Schwarz and Clore (1983) to affect how people view their happiness and life satisfaction. When participants in the study were made aware of the mood primer, which was the weather in their experiment 2, the effect of the primer disappeared only for those in a bad mood. This study was among the first ones to lay ground for the theory of feelings-as-information, which in general postulates that "people attend to their feelings as a source of information" (Schwarz, 2011, p. 294). The decision making at the date of the sale may thus be affected by the sunlight exposure on that day. However, this must be seen in light of the fact that making offers to a seller is not a situation characterized as a routine for inexperienced buyers. As summarized by Schwarz and Clore (2007), there is a distinction in the cognitive process when facing a problematic and a benign situation. When facing a problematic situation, the mind takes an analytical approach driven by the need to act, while facing a benign situation does not require this thus it relies on previously learned routines (Schwarz & Clore, 2007, pp. 408–410). To this extent, demanding processing of difficult topics, with potentially large amounts of information, are theorized to be "infused" by affect (Forgas, 1995).

There two main channels that may be important for the role of the mood induced by sunlight at the showings. First, when deciding about making offers to the seller, this requires the buyer to recall the experience from the showing and all other information retrieved. If the buyer had her mood influenced by the sunlight exposure at the showing, then memories about different characteristics of the housing unit is loaded with a particular mood (Bower, 1981). In a sense, this is related to the mood priming used by Schwarz and Clore (1983), and it is directly relevant for the situation at hand. Second, the mood of the buyer at the date of the sale is relevant for which memories that are most available to her, which is called mood-congruent recall (Bower, 1981; Clore et al., 1994; Forgas, 2017). Specifically, if the buyer is in a good mood, memories with mood that is consistent with the current mood will become more available.[5] Hence, if most of the experience from the showing was loaded with bad mood due to little sunlight, and if the buyer is already in a bad mood, this may amplify a bias from misattribution of mood in decision making on the date of the sale.

Finally, in line with a general feelings-as-information theory, while not in line with the mood-congruent recall theory, Johnson and Tversky (1983) find that when inducing participants in an experiment with information about a bad event, this increased their expectations about likelihoods of bad events in general, and not only specifically related to the primer. Put differently, with a difficult task at hand, people simplify the case by attributing information from their feelings, independently of the similarities in content (Schwarz & Clore, 2007, p. 395).

Along with the prevalence of seasonal depression, especially in the fall and winter, as well as the stress and uncertainties associated with buying houses, buyers may be in a bad mood. In turn, this could make buyers recall not only the actual bad aspects of the house, but also their memories from the showings that were loaded with bad mood. Depending on the latitude at which buyers attend showings, this may make them to be more affected by the lack of exposure to sunlight during the darker periods of the annual cycle.

1.2.4 Search in the housing market

The housing market search literature mostly consists of theoretical studies. The purpose of applying search theory to the housing market is to help understand certain aspects of the transaction process between buyers and sellers, specifically how they search and match. The key concept in search models is search frictions, which are used to explain market outcomes. Examining the role of search frictions in transaction processes helps understanding how the housing market, and specifically sub-markets, clear (Han & Strange, 2015). There are two main types of search model: models with random search and models with directed search. In what follows, these types are consider in a housing market context.

Studies with random search let the matching of sellers and buyers be random. The rate at which they meet is determined by a matching function which depends on the number of buyers and sellers in the market, for instance as presented in the baseline model of Genesove and Han (2012). This model predicts that a positive demand shock leads to

[5]This is related to the availability heuristic presented by Tversky and Kahneman (1973).

an increase in the acceptance rate, thus buyers will visit fewer houses before buying (matching), and a decrease in sellers' time-on-market. How this affects the probability of purchase of the buyers is unclear. As pointed out by Han and Strange (2015), the baseline model of Genesove and Han (2012) can explain the positive correlation between prices and volume in the housing market.[6]

While some studies keep the buyers-to-sellers ratio constant, some allow it to vary over time. An example is Novy-Marx (2009) who allows market entry of both sellers and buyers to depend on the value of doing so. The importance of endogenous entry is made clear by what he calls a "self-reinforcing feedback loop" (Novy-Marx, 2009, p. 1): when an exogenous shock improves the market conditions in favor of the sellers, such as an inflow of more buyers, this will reduce the time-on-market leading to a larger buyers-to-sellers ratio and tighter market, which again further improves the market in favor of the sellers. In total, the model is able to explain how search frictions can lead to high volatility in the market.

An important aspect of the housing market is that buyers often are sellers in the same sub-market and vice versa. This makes the housing market differ from, for instance, the labor market, which is the typical market to apply search models to.[7] This aspect is included in the random search model presented by Wheaton (1990), in which sellers buy first thus owning two housing units at the same time, and then they wait until a buyer is matched to their old unit. Importantly, the model defines three occupancy types of households: households living in a matched house type; households living in a matched house type while also owning the old mismatched house type; and households living in a mismatched household. The incentive for households to start searching comes from the change in the occupancy type from matched to mismatched, and this is induced by changes in aggregate supply and demand. Thus, the model incorporates both random search and turnover, and it explains the relationship between vacancy, turnover, and prices. Specifically, the model predicts that high turnover can lead to higher prices, and the existence of a structural vacancy (Wheaton, 1990).

The intuition about turnover is expanded upon by Ngai and Sheedy (2020), who model the decision to relocate is as being endogenous. The homeowners keep their home due to the match quality being above a certain threshold. The match quality is however prone to idiosyncratic changes, which are caused by shocks represented by life events (Ngai & Sheedy, 2020, p. 2489). Moreover, the model accommodates for the importance of

[6]Recall that loss aversion could explain the positive correlation between prices and volume in the housing market. This was a motivating factor for Genesove and Mayer (2001) but also for their prior study of equity constraints (Genesove & Mayer, 1997).

[7]See e.g. the text book of Pissarides (2000). Note that the importance of buying before selling in the housing market is investigated by Moen et al. (2021) using a random search setup.

shocks in the macroeconomy and policies by letting these affect the threshold for deciding to move. This model is used to explore the importance of housing misallocation and relocation, suggesting that macro-shocks can lead to the lowest match quality homeowners to relocate, which they call a "cleansing effect". Moreover, when a cleansing occurs less homeowners are close to the relocation threshold, thus resulting in a new long run steady state (Ngai & Sheedy, 2020, p. 2490).

Studies with directed search let the matching of sellers and buyers be directed by list prices. As noted by Han and Strange (2015), in general, list prices can affect selling prices in two ways, namely in the bargaining between sellers and buyers (Yavas & Yang, 1995) and to attract more buyers (Horowitz, 1986). List prices can serve different roles for the sellers. First, they can serve as commitments, in which list prices work as constraints on selling prices that these sellers will accept. Meaning, list prices work as ceilings on selling prices. Such a role is used in the model of Y. Chen and Rosenthal (1996), who find that lower list prices should encourage more buyer search.[8] Moreover, in a monopolist model, they find that when this seller has lower bargaining power, the list price becomes less important.

Second, list prices can serve as partial commitments. This is more closely related to how list prices seem to work in the past decades, namely that housing units are sometimes sold for more than their list prices. This is both empirically documented and theoretically addressed by Han and Strange (2016). Their model predicts that reducing the list price is only a feasible strategy to attract buyers until a certain point, and whether the selling price becomes higher than the list price depends on the number of buyers that have been attracted and their match quality. Importantly, when there is at least two buyers this leads to an auction, and with the possibility of at least two buyers having match values above the list price, then there is a non-zero probability of the selling price being larger than the list price.

Third, list prices can be instruments for signaling seller type. Potentially the most direct implementation of this is the model of Albrecht et al. (2016), who separate sellers into two types based on their reservation price. A seller with a higher reservation price is "relaxed" and a seller with a lower reservation price is "motivated". They find an efficient separating equilibrium where the relaxed sellers put higher list prices and the motivated sellers put lower list prices, in which reducing list prices lead to more buyers.

Empirical studies of search are rarer than theoretical ones. Some studies take an empirical approach to test their theoretical predictions, such as Carrillo (2012), Han and

[8]Y. Chen and Rosenthal (1996) distinguish between the terms *list price* and *ask price*, in which the former means the list price on the market, while the latter means that this is not a fixed price but more of a suggested price potentially dependent on the seller's reservation price. List prices in housing markets, such as the Norwegian housing market that is analyzed in the thesis, are ask prices.

Strange (2016), Haurin et al. (2010), Moen et al. (2021), and Rekkas et al. (2022), among others. Other studies supplement with analyzing survey data (Genesove & Han, 2012; Han & Strange, 2016). Larger datasets of search, and not just prices, are rare, but there are some studies that utilize email alert data of new listings (Piazzesi et al., 2020) and data on advertisement clicks (Barnwell & Fournel, 2022). Still, the general lack of search data is prominent in the literature, which could be related to the focus on list prices.[9]

1.3 Data

This section describes the data analyzed in the thesis, including descriptions of important institutional factors.

1.3.1 Data sources

There are two main datasets analyzed in this thesis, and both are proprietary. First, a dataset comprising on-market transactions between individuals from the most populated municipalities in Norway. These data are provided by Eiendomsverdi AS, a private firm that daily collects information about housing transactions in Norway from the Land Registry and real estate agencies. This dataset includes list and selling prices, appraisal values, time of listing and sale, attributes, zip codes, and consistent identifiers for the housing units, buyers, and sellers. These data are joined with information about the buyers and sellers provided by Statistics Norway. Information on individuals include income, wealth, debt, gender, age, education level, occupation category, household composition, and a household identifier. The data on individuals have a panel structure.

Second, a dataset comprising on-market transactions from the second largest real estate agency in Norway, DNB Eiendom. This is a rare dataset in the sense that it includes information at the bid level, with identifiers of bidders that are consistent across transactions for parts of the available time period. In addition to list and selling prices, attributes, and addresses, the dataset contains information about the time that bids are received by the real estate agent (hereafter realtor), when they expire, at which time they are accepted, and the size of the bids. Together with the bidder identifiers that are always consistent within transactions, this makes it possible to derive additional details about the bidding process, such as its duration, and how many bidders there are. In the subsample with consistent bidder identifiers both within and between transactions, it is possible to observe bidders over time. Moreover, the data include search information: how many showing participants there have been at the public showings; the dates of the first and the last public showings; how many showings there have been in total; how many

[9]There are other studies not founded in search theory about the role of list prices for selling prices and time-on-market, such as Andersen et al. (2022), Anglin et al. (2003), and Anundsen et al. (2022).

people have signed up to receive information about incoming bids; and the total number of advertisement clicks. There is also a smaller subsample of more recent transactions that includes information about all public showings. These include the exact times at which the showings were conducted and the number of participants.

The second and third paper also utilize weather data from the Copernicus Climate Change Service (Copernicus Climate Change Service, 2024a, 2024b, 2024c). These include both in-situ, satellite, and forecasted variables of different weather outcomes, such as cloud cover, temperature, and precipitation.[10]

1.3.2 Institutional context

The data analyzed in the thesis are from the Norwegian housing market. When a homeowner wants to sell her housing unit in this market, the homeowner hires a realtor to help sell the unit. The selling process is a multi-stage process that includes gathering information about the unit; acquiring estimates of the market value; creating the prospectus and advertisement material; listing the unit on the market in order to attract buyers; conduct public and private showings; keep track of incoming offers; and facilitate for finalizing the contract of the transaction between the final buyer and the seller. Although the realtor is hired and paid by the seller, the realtor is mediator who by law also must care for the interest of the buyers (Eiendomsmeglingsloven, 2007, §6-3).

There are some specifics about the selling process that are relevant for the thesis. First, in the information collection stage, the realtor hires a professional surveyor, a specialist appraiser, who inspects the housing unit and provides a technical report of the findings. The report is usually provided as an attachment in the prospectus. In some municipalities, surveyors also provided their own estimate of the market value, namely an appraisal of the value. The appraisal value was used as an indicator for the realtor and the seller in choosing the list price, and it was made available to prospective buyers. The practice of using appraisal values ended in mid-2016 when the Norwegian Real Estate Association decided to switch to predictions of market values generated by an Automated Valuation Model (AVM).[11] These specific AVM predictions are not available to prospective buyers.

Second, the listing of housing units are primarily online, at the internet marketplace Finn.no, but some also advertise in newspapers. Newspaper advertisements have become rare in recent years. Therefore, when it comes to attracting buyers, the online advertisement at Finn.no is more or less the sole channel to do so. This online listing can be updated by the realtor, for instance with the public showing schedule. Showings are

[10] In-situ data are collected at the location the weather occurs, contrasting the data collection from a distance called remote sensing. Observations from satellites are remote sensing.

[11] One of the data providers in this thesis, Eiendomsverdi AS, is the owner, developer, and maintainer of this AVM. The AVM is at the core of their business model.

usually conducted by the realtor or a colleague from the realtor office, and during these showings, people can sign up for being informed about incoming offers.

Third, as noted in the introduction, the stage in which buyers make offers is resembling an open ascending bid auction. The bidding process differs in some ways from a typical auction, such as the fact that the winning bidder is the one whose bid is accepted by the seller. The winning bid does not need to be the highest bid. Even when bids have been received, this does not mean that the unit will be sold, because the sale is finalized only when a bid is accepted. When there is a single bidder, the process of placing bids and receiving counteroffers from the seller becomes a negotiation. In negotiations without counteroffers, the recorded negotiation in the data resembles a bidder that has competed against herself. Despite some differences from a more standardized auction, the third and fourth papers refer to the bidding process as being an auction.

In addition, all papers in the thesis take in some way account of the different housing ownership types in Norway. There are two main ownership types, housing cooperatives (hereafter co-ops) and non-co-ops.[12] An owner of a co-op apartment does not own it directly, but has a share in the housing cooperative that gives the right to access and use that particular apartment. Co-ops are not exclusively apartments, for example, some row houses have this ownership structure, but they are very rare. Non-co-ops are either apartments, detached houses, semi-detached houses, or row houses. These are owned directly by the owner.[13] This also goes for block apartments. When the ownership type is referred to as owner-occupier in the fourth paper, this refers to non-co-op ownership, and these can in principle be renter-occupiers of non-co-ops.

1.4 Methodological concepts

Common for all papers in the thesis is that their contributions are in the form of empirical evidence. The first paper is directly dealing with issues of identification within an established identification framework. The second paper uses a hedonic model, with home characteristics, high dimensional fixed effects, and de-seasoned control variables of the weather as the empirical strategy. The third paper employs machine learning to isolate the importance of list prices, drawing intuition from a search theoretical perspective. The fourth paper uses the same identification strategy as the second paper, namely high dimensional fixed effects, but in a different setting. While these are the identification designs used in the thesis, this section describes important specifics of the main concepts

[12]There exists other types but they are rare. For several reasons, including poor data quality and different mortgage conditions, these types are omitted when analyzing the data. The exception is the third paper, which keeps the other ownership types in the sample because these observations may provide relevant information to the machine learning algorithm.

[13]Ownerships of both co-ops and non-co-ops are registered in the Land Registry.

relevant for the employed methodological approaches.

1.4.1 Implicit prices, omitted variable bias, and measurement error

The three papers that employ regression for estimating econometric models are doing so to identify different effects. The dependent variable is in most cases a price, being either the list price or the selling price. When estimating a relationship, one must also account for the information that is relevant for the dependent variable and that correlates with the variables of interest, meaning the explanatory variables. If one is not able to properly control for omitted variables, this leads to omitted variable bias, resulting in estimates that are either too small or too large depending on the direction of the bias. Omitted variable bias is a key concern in the thesis.

The housing market is highly heterogeneous, and the uniqueness of housing units makes no two identical units exist. This even goes for block apartments because they have slightly different view even when situated side-by-side with identical layout. Still, similar units should have similar prices, which relates to a widely used concept in the housing market literature, namely hedonic prices. The theory of hedonic prices is rooted in consumer theory, with the idea being that consumers have preferences about the characteristics of goods, and the proportion of these, rather than just the goods themselves (Lancaster, 1966). Moreover, because consumers care about the characteristics, the market prices of goods should depend on the characteristics, implying the existence of a market for implicit prices. The theory of implicit (hedonic) prices is famously explored and formalized by Rosen (1974).

Depending on the research question, implicit prices and hedonic predictions generated by hedonic models are used in different ways, sometimes to accommodate concerns about omitted variable bias, and sometimes to provide a measure of market prices. For instance, the first paper follows the literature by using hedonic predictions of selling prices as an explanatory variable. Doing so makes it possible to investigate whether the explanatory variables of interest can explain the difference between the list price and the predicted price. In that paper, the hedonic prediction of a house price is supposed to capture the seller's actual expectation of the market price, while in other cases the prediction is meant as a general market price estimate (e.g., Anundsen et al. (2022)).

In the second paper, hedonic characteristics are used as control variables to reduce potential omitted variable bias. These variables are important for capturing the segmentation effects that can be picked up by the explanatory variables of interest. Predictions of selling prices that are produced by a hedonic model that accounts for important market segments will nest these effects, but it is not sufficient for the purpose of this paper to only use a compound of all these variables. Although a compound price puts weights on each of the predictors, a prediction is still just the sum of the housing characteristics weighted by their estimated implicit prices. This is a less flexible solution than including the characteristics themselves as control variables. The importance of the characteristics are thus not fixed by the first stage hedonic model estimation, and including them as controls allows their them to pick up other potential effects when including the explanatory variables of interest. Implying, the more suitable approach depends on the research question.

An important concern in housing market research and hedonic modeling is heterogeneity unobserved by the econometrician. When training a model with the purpose of predicting selling prices, unobserved heterogeneity consists of everything that is important for selling prices but are not available in the data. Unobserved heterogeneity in selling prices is not a concern for identification when estimating a selling price model if what is unobserved is unrelated to the explanatory variables. It only becomes a problem if the missing information is correlated with the explanatory variables of interest. As is clear, unobserved heterogeneity can lead to omitted variable bias.

Solutions for dealing with issues related to unobserved heterogeneity are depending on the type of heterogeneity that causes problems. There are two types of heterogeneity: time-invariant and time-varying heterogeneity. Bias that occurs from omitting the former can be mitigated by utilizing repeat sales data, but this may not always be a feasible solution. The latter can be mitigated by collecting more information about the housing units and the location that captures changes in relevant factors over time, but this is typically unobtainable information. In general, there are no simple solutions for dealing with time-varying unobserved heterogeneity.[14]

A related but different issue is measurement error. The datasets analyzed in the thesis are register data, and there are both missing values and clerical errors in them. Yet, these errors are most likely, and therefore assumed to be, random. When explanatory variables are measured with some error, this leads to attenuation bias. Still, in the first paper, the key problem is related to when measurement error is not random. In particular, the measurement error is important for the dependent variable and other explanatory variables. Meaning, the measurement error itself is an omitted variable. This may arise when a variable is measured by a prediction of itself, and when the unobserved part, being the prediction error, is correlated with the dependent and explanatory variables. Such a measurement error is referred to as differential Berkson error in the first paper (Berkson, 1950).[15]

[14]Still, as employed in the first paper in thesis, there are some weaker solutions that may help controlling for time-varying unobserved heterogeneity when estimating the economic relationships.

[15]Note that differential measurement error has gotten some recent attention for its role in house price indices (Hayunga et al., 2024). It is also addressed in the estimation of demand concavity (Guren, 2018).

1.4.2 Machine learning with boosted trees

The third paper uses machine learning to model buyer arrival to listed housing units. While linear regression estimates the linear relationships between the dependent and the explanatory variables, methods of machine learning allows for highly non-linear relationships.[16] Together with the fact that machine learning methods are learning from data, in some cases iteratively, in order to generalize out-of-sample, they provide the flexibility required to perform many prediction tasks.

The machine learning algorithm encountered in the third paper is called eXtreme Gradient Boosting (hereafter XGBoost) (T. Chen & Guestrin, 2016). XGBoost builds an ensemble of decision trees through a sequence of iterations, in which the tree in the next iteration is the one that optimizes a chosen objective function. This procedure takes into account what is already learned from the data, so that each new iteration adds a tree for improving the current model.[17] In this sense, XGBoost is an automated model selection algorithm which depends on a set of hyper-parameters used to control the learning procedure. For prediction purposes, using decision trees implies taking account of interactions between predictors. This is an advantage when the outcome to predict is the result of many different variables interacted with each other. While XGBoost creates a highly complex sequence of decision trees that is hard to interpret directly, explanatory tools exist for this purpose. A commonly used tool is SHAP values as presented by Lundberg and Lee (2017). SHAP values represent the contribution to the predictions from each of the predictors, which is provided at the level of the unit of observation. The average of the absolute SHAP values will thus indicate the importance of a predictor for the model. In total, these values reveal how a trained model is treating the different predictors.

1.5 Synthesis of papers

As described, the thesis considers different aspects of seller and buyer behavior in the housing market, and as such, they are related in this dimension. The thesis contributes in general to the housing market literature, yet, some contributions goes beyond this field. Specifically, the first paper shows the importance of keeping track, and the implications, of differential measurement error, while also arguing that loss averse behavior can be institutionally dependent. The second paper contributes with insights of the effect of sunlight on housing prices. The third paper offers new insights into buyer search, and the

[16] Some methods, for instance support vector machines, models relationships based on fitting multi-dimensional hyperplanes (Lantz, 2023, p. 265).

[17] More information including technical details are provided by T. Chen and Guestrin (2016) and on the XGBoost website https://xgboost.readthedocs.io/en/stable/.

fourth paper provides the first empirical evidence of the importance of bidding experience for adjustments in subsequent bidding behavior.

1.5.1 On Loss Aversion in Housing Markets: Evidence from Norway

The topic of seller behavior consistent with loss aversion in the housing market has been the subject of several studies, starting with Genesove and Mayer (2001) and followed by Andersen et al. (2022), Anenberg (2011), Bokhari and Geltner (2011), Einiö et al. (2008), Engelhardt (2003), Lamorgese and Pellegrino (2022), and Zhou et al. (2022), among others. Specific for the literature is the hardships of identification, in which there are known issues related to the inability to obtain measurements of the sellers' true expectations of selling prices when they list their housing units. Still, the existing literature finds evidence which is attributed to loss aversion. This paper builds on the existing literature by focusing on the seminal work of Genesove and Mayer (2001) and proposing a different measurement of sellers' price expectations.

The proposed measurement is appraisal values provided by surveyors who visit the housing units and assess both their technical properties and potential market value. The appraisal values should nest more of the important information about the housing units and they serve as supportive tools in list price decisions. From an econometric perspective appraisal values will help both explaining the dependent variable, that being list prices, and reduce biases due to omitted quality that arise when measuring expectations with hedonic predictions, which is the usual way to do this. Yet, appraisal values have also an institutional aspect to them in the Norwegian housing market: they are observed by buyers and thus are usually upper bounds for list prices.

Transaction data from the Oslo housing market are used to construct a dataset of purchases-to-listings in which the buyers at the time of purchase are required to perfectly match the sellers at the time of listing. Utilizing these data to estimate the loss aversion effect produces two different results. First, replication gave results consistent with other studies (Bokhari & Geltner, 2011; Genesove & Mayer, 2001). Second, using appraisal values as measurements of sellers' price expectations gave very small and insignificant estimates of the effect. In light of the benefits of appraisal values relative to hedonic predictions, the latter results should be preferred. A step-wise examination of the difference between the two results, that addresses the estimation biases due to unobserved time-invariant and time-varying heterogeneity, reveals that the replication results suffer from omitted variable bias to a larger extent than previously believed.

The findings provide insight into the importance of the institutional context for acting in accordance with behavioral impulses. Importantly, loss aversion has been argued to explain some stylized facts of the housing market, such the positive correlation between prices and volume especially in market busts. Because this pattern is prominent in the Norwegian housing market[18] and loss averse seller behavior is absent, this requires alternative explanations. Notably, some theoretical search models are consistent with the relationship between prices and volume (Anenberg, 2016; Genesove & Han, 2012).

1.5.2 Here Comes the Sun: The Effect of Sunshine on Home Prices

The importance of sunshine, and the lack thereof, on asset prices in financial markets has been the topic of several studies, such as Goetzmann and Zhu (2005), Goetzmann et al. (2015), Hirshleifer and Shumway (2003), Kamstra et al. (2003), and Saunders (1993). However, because most households do not invest a large share of their wealth in the stock market, but rather in the housing market, the insights from these studies are of limited application for household finance. The study that comes closest to the second paper in the thesis is Gourley (2021) who investigates the effect of temperature and precipitation on house prices in Jefferson County, Colorado. While Gourley finds evidence suggesting that both weather measures affect prices, he remains agnostic to which points in time during the transaction process the weather affects price formation.

This paper investigates the effect of cloud cover on house price formation, specifically by isolating the effect from two important stages in the transaction process where buyers form beliefs and make decisions: at public showings and when placing bids. Separating the effect from showings and the day of the sale is made possible by utilizing several datasets from the Norwegian housing market, a market characterized by low time-on-market and standardized on-market transaction processes. The main dataset comprises on-market transactions with information about selling prices, list prices, typical housing characteristics, and consistent identifiers of housing units and buyers. These data are combined with information about the buyers to investigate heterogeneous effects across different buyer segments. Realizations of daily cloud cover are measured using satellite observations, which are aggregated to the zip code level. In addition, the analysis is supplemented with a dataset of bids from on-market transactions, which includes information about the number of bidders in transactions, showing participation, and showing dates.

The results suggest that prices are affected by sunlight, measured by cloud cover, at the showings but not on the day of the sale. The effect is found to be prominent when the days are shorter and darker, but not during the summer. Importantly, the paper does not find that cloud cover affects factors related to market liquidity. Together, this is interpreted as suggestive evidence for the effect being due to misattribution of mood, and not due to sunlight affecting factors related to market liquidity. A plausible channel

[18] This pattern can be seen, for example, at the webpage of Real Estate Norway, the association for realtor brokerage firms in Norway, https://eiendomnorge.no/boligprisstatistikk/statistikkbank/. The webpage is in Norwegian.

for how sunlight at the showings affects prices is through sunlight-induced mood encoded into memories. Another explanation of the effect is that sunlight enhances the perceived value of features of the housing unit, and non-apartments should benefit more from this. On the contrary, there is only a significant sunlight effect on apartments. Furthermore, the results suggest that those buying alone, buyers with lower gross wealth, and first-time buyers are affected by sunlight. This supports the proposed mood channel: being in a position with low bargaining power or purchasing power may lead to more stress and uncertainty. In turn, this can make them more inclined to rely on mood as a source of information, either directly from their feelings (Schwarz & Clore, 1983) or amplified through mood-congruence (Bower, 1981).

1.5.3 How directed is housing search?

The housing market search literature is primarily focused around models in which the matching of buyers and sellers is either random or directed by list prices. These represent the two main approaches of modeling search. Data on buyer search is hard to come by, but some studies, including Barnwell and Fournel (2022) and Piazzesi et al. (2020), have made use of data specifically on the initial stage of information collection, which Han and Strange (2015) refer to as the "pre-search" stage. Still, empirical evidence of buyer search remains scarce in the literature.

This paper provides evidence of the importance of list prices for buyer search in the housing market. From a directed search perspective, list prices should be important for search because they provide relevant information about the sellers. For example, sellers may use list prices as commitments and to signal their types. The analysis utilizes buyer search outcomes as represented by buyer arrival to listed housing units in the Oslo housing market. In the data, buyer arrival consists of internet advertisement clicks, showing participation, the number of people registered for receiving information about new bids, and the number bidders. The paper acknowledges that buyer arrival may result from highly complex processes, potentially to a greater extent than what has been studied previously. To alleviate concerns about modeling buyer arrival that is parametrically constrained, buyer arrival is modeled using a decision tree ensemble method called XGBoost.

The machine learning algorithm is provided a large set of predictors, including hedonic characteristics, location, time of sale, local market factors, a market tightness measure, distances to amenities, the weather, and previous arrival stages. Then the importance of the list price for arrival is assessed by training models both with and without the list price and comparing out-of-sample prediction performance. The findings suggest that list prices have in general little to say for buyer arrival when already accounting for much other relevant information that buyers may have preferences for. Yet, for cheaper housing units, that being units with list prices below approximately 400,000-450,000 USD, list prices affect advertisement clicks and the number of bidders in a manner consistent with directed search.

1.5.4 Repeat bidder behavior in housing auctions

Despite being concerned about matching and mismatching in the housing market, the search literature has not focused its attention on how buyers change their behavior when failing to win in the bidding processes for housing units (hereafter auctions). In fact, there is a lack of literature about how agents in the housing market behave over time, being either buyers, sellers, or realtors. In a recent study, Gilbukh and Goldsmith-Pinkham (2023) investigate the importance of realtor experience, finding that hiring inexperienced realtors to sell housing units lead to lower probability of sale. This study comes the closest to the fourth paper in the thesis which is concerned with how buyer experience affects buyer behavior, thus being related in one dimension: the importance of housing market experience of market agents.

This paper is the first study to document bidder behavior as they move across auctions and gain experience. Bidders accumulate experience by bidding in, and losing, housing auctions. The analysis is made possible by a dataset of bids from housing transactions in Norway which includes consistent identifiers of bidders over time. Because the dataset includes information about bids and transactions, the bids of the bidders are compared with their auction competitors and transaction-specific factors, such as list prices. In total, the paper estimates the effects of experience on repeat bidder behavior, in which it distinguishes between those who are observed to win (hereafter winners) and those who are not (hereafter losers).

The findings suggest that bidders change their behavior when they lose in housing auctions, and that there are differences between winners and losers. In total, the results indicate that winners are changing their behavior in favor of increasing their likelihood of winning, and to a larger degree than losers. The winners adjust their search towards smaller and cheaper houses. While being constrained by their budgets, winners increase their competitiveness even when the size of their bids are unchanged.

These findings are also relevant for the search literature. The effects are estimated while controlling for bidder fixed effects, so that some factors related to match quality are accounted for. Because winners adjust their search towards lower priced housing units, this implies that, provided that one controls for bidder specific attributes, list prices matter across the whole list price distribution. Hence, this goes beyond the initial search effort when entering the market. Yet, although contrasting the findings in the third paper, this result is still prone to be driven by the fact that list prices nest much information about housing unit characteristics, location, and neighborhood. Thus, these findings are not evidence consistent with list prices being important instruments for sellers to provide information about themselves to attract buyers.

1.6 Limitations and further research

The papers in the thesis have several limitations. The paper about loss aversion attempts to deal with omitted variable bias in order to provide credible estimates of sellers' loss averse behavior. Using appraisal values has one key drawback. Appraisal values will nest more of the important information leading to estimation bias, but appraisal values are made available to prospective buyers. This visibility poses a threat to identification, because it makes sellers less likely to display their loss aversion if they truly are loss averse. Meaning, the sellers may feel constrained in their list price choices, making them unable or unwilling to behave in accordance with being loss averse. However, the paper does not attempt to test the hypothesis of constrained sellers, but a natural cut-off that future research should consider is the switch in mid-2016 from using appraisal values to AVM predictions as tools for evaluating the potential market values. This requires getting hold of the actual AVM predictions supplied to the sellers, but these are not available for researchers.

There are also other ways, besides hedonic predictions of selling prices and appraisal values, to measure these expectations. The paper highlights the importance of being able to measure expectations with the least measurement error as possible. This is the key for identification. If price expectation measurements are prone to measurement error, unbiased but still potentially attenuated estimates can be found if the measurement error is not related to unobserved quality of the housing units. A potential way of measuring expectations is to ask sellers what they expect to get for their house shortly before deciding upon the list price. Still, a different threat to identification is the fact that the model framework relies on the assumption that the purchasing price is the reference point. Some might argue, in line with Kőszegi and Rabin (2006), that the reference point could be the seller's expectation "...held in the recent past" (Kőszegi & Rabin, 2006, p. 1141). Meaning, future research should also be concerned about measuring the relevant reference point, which may be an expectation itself.

The paper about the effect of sunlight, measured by cloud cover, on prices utilizes two transaction datasets with different advantages and limitations. The main dataset does not include information about the exact showing dates while this is found in the supportive dataset. However, the supportive dataset is not suited for the analysis because it is too small and consist of transactions from a single realtor agency. This makes it unfitted for controlling for unit fixed effects which are crucial for the identification strategy. Having actual showing dates in the main dataset would lead to less noise in the cloud cover measurements. The sample size is also limited, and a longer time period should allow for more repeat sales. The study may also suffer from the lack of high definition, local, cloud cover data at the precise intra-day showing times. Also, even unit fixed effects does not deal with time-varying heterogeneity, which in the context of the paper relates to how the potential for sunlight exposure varies with factors such as urban densification.

The paper about whether buyer search is directed by list prices also has some shortcomings. Buyer information is not taken into account. Budgets and preferences in the time period when buyers are searching, which are related to household constellation, income, wealth, and employment, can help capturing the link between these buyers and the listed housing units. This link is the idiosyncratic match quality. Although being a different perspective, this might alleviate some uncertainties about market segments in which the buyers are searching. Meaning, list prices may be more important within market segments. Such an exercise may require a larger sample. This leads to another limitation of the paper, namely the sample size, which may be limiting the potential for machine learning to fully learn from the data because there are too few cases to learn from.

Finally, the results presented in the paper about repeat bidders are limited by the fact that the data are from a single realtor agency. The agency is the second largest in Norway, but many of the bidders observed repeatedly are likely bidding on housing units at other agencies. There are three main implications. First, bidders observed in the data are likely also bidding on other unobserved housing units in the sample period. Because bidders should only care about the characteristics of housing units on the market independent on the realtor agencies, sample selection should not be an issue. Yet, the estimates depend on the frequency of which bidders are observed, thus not precisely capturing changes in bidding behavior. Second, the analysis splits bidders into winners and losers, but those not observed winning are still likely to win at some point in time, probably at another realtor agency. Third, related to the first two points, is that bidders are observed at different stages in their search process. The first time a bidder is observed does not mean that the bidder has not been active in other auctions earlier. This bidder may potentially be in the late stage of her search process, and she may have already adjusted her bidding behavior based on unobserved experience. Thus, a more complete dataset may be required to make more viable conclusions.

References

Albrecht, J., Gautier, P. A., & Vroman, S. (2016). Directed search in the housing market. *Review of Economic Dynamics*, *19*, 218–231.

Andersen, S., Badarinza, C., Liu, L., Marx, J., & Ramadorai, T. (2022). Reference Dependence in the Housing Market. *American Economic Review*, *112*(10), 3398–3440.

Anenberg, E. (2011). Loss aversion, equity constraints and seller behavior in the real estate market. *Regional Science and Urban Economics*, *41*(1), 67–76.

Anenberg, E. (2016). Information frictions and housing market dynamics. *International Economic Review*, *57*(4), 1449–1479.

Anglin, P. M., Rutherford, R., & Springer, T. M. (2003). The trade-off between the selling price of residential properties and time-on-the-market: The impact of price setting. *The Journal of Real Estate Finance and Economics*, *26*, 95–111.

Anundsen, A. K., Nenov, P., Larsen, E. R., & Sommervoll, D. E. (2022). *Pricing and incentives in the housing market* (Housing Lab Working Paper No. 3). Oslo Metropolitan University.

Barberis, N. C. (2013). Thirty years of prospect theory in economics: A review and assessment. *Journal of Economic Perspectives*, *27*(1), 173–196.

Barnwell, J.-L., & Fournel, J.-F. (2022, April). *Seller's (Mis)Fortune in the Housing Market: Directed Search in Online Real Estate Platforms* (Working paper). Retrieved October 24, 2024, from https://www.jeanfrancoisfournel.com/uploads/directed_search.pdf

Beggs, A., & Graddy, K. (2009). Anchoring effects: Evidence from art auctions. *American Economic Review*, *99*(3), 1027–39.

Berkson, J. (1950). Are there two regressions? *Journal of the American Statistical Association*, *45*(250), 164–180.

Bokhari, S., & Geltner, D. (2011). Loss aversion and anchoring in commercial real estate pricing: Empirical evidence and price index implications. *Real Estate Economics*, *39*(4), 635–670.

Bower, G. H. (1981). Mood and memory. *American psychologist*, *36*(2), 129.

Brown, J., & Matsa, D. A. (2020). Locked in by leverage: Job search during the housing crisis. *Journal of Financial Economics*, *136*(3), 623–648.

Campbell, J. Y., & Cocco, J. F. (2007). How do house prices affect consumption? evidence from micro data. *Journal of Monetary Economics*, *54*(3), 591–621.

Carrillo, P. E. (2012). An empirical stationary equilibrium search model of the housing market. *International Economic Review*, *53*(1), 203–234.

Chen, T., & Guestrin, C. (2016). Xgboost: A scalable tree boosting system. *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, 785–794.

Chen, Y., & Rosenthal, R. W. (1996). Asking prices as commitment devices. *International Economic Review*, 129–155.

Clore, G., Schwarz, N., & Conway, M. (1994). Affective causes and consequences of social information processing. In R. S. Wyer Jr & T. K. Srull (Eds.), *Handbook of social cognition: Volume 1: Basic processes* (pp. 323–417, Vol. 1). Psychology Press.

Copernicus Climate Change Service. (2024a). Cloud properties global gridded monthly and daily data from 1982 to present derived from satellite observations. Copernicus Climate Change Service (C3S) Climate Data Store (CDS). https://cds.climate.copernicus.eu/datasets/satellite-cloud-properties?tab=overview

Copernicus Climate Change Service. (2024b). ERA5-Land hourly data from 1950 to present. Copernicus Climate Change Service (C3S) Climate Data Store (CDS). https://cds.climate.copernicus.eu/cdsapp#!/dataset/reanalysis-era5-land?tab=overview

Copernicus Climate Change Service. (2024c). Nordic gridded temperature and precipitation data from 1971 to present derived from in-situ observations. Copernicus Climate Change Service (C3S) Climate Data Store (CDS). https://cds.climate.copernicus.eu/cdsapp#!/dataset/insitu-gridded-observations-nordic?tab=overview

Eiendomsmeglingsloven. (2007). Lov om eiendomsmegling (LOV-2007-06-29-73). https://lovdata.no/dokument/NL/lov/2007-06-29-73

Einiö, M., Kaustia, M., & Puttonen, V. (2008). Price setting and the reluctance to realize losses in apartment markets. *Journal of Economic Psychology*, *29*(1), 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*(1), 171–195.

Forgas, J. P. (1995). Mood and judgment: The affect infusion model (aim). *Psychological bulletin*, *117*(1), 39.

Forgas, J. P. (2017). Mood effects on cognition: Affective influences on the content and process of information processing and behavior. In M. Jeon (Ed.), *Emotions and affect in human factors and human-computer interaction* (pp. 89–122). Elsevier.

Garrett, I., Kamstra, M. J., & Kramer, L. A. (2005). Winter blues and time variation in the price of risk. *Journal of Empirical Finance*, *12*(2), 291–316.

Genesove, D., & Han, L. (2012). Search and matching in the housing market. *Journal of Urban Economics*, *72*(1), 31–45.

Genesove, D., & Mayer, C. (2001). Loss Aversion and Seller Behavior: Evidence from the Housing Market. *The Quarterly Journal of Economics*, *116*(4), 1233–1260.

Genesove, D., & Mayer, C. J. (1997). Equity and time to sale in the real estate market. *The American Economic Review*, *87*(3), 255.

Gilbukh, S., & Goldsmith-Pinkham, P. (2023, September). *Heterogeneous real estate agents and the housing cycle* (Working Paper No. 31683). National Bureau of Economic Research.

Goetzmann, W. N., Kim, D., Kumar, A., & Wang, Q. (2015). Weather-induced mood, institutional investors, and stock returns. *The Review of Financial Studies*, *28*(1), 73–111.

Goetzmann, W. N., & Zhu, N. (2005). Rain or shine: Where is the weather effect? *European Financial Management*, *11*(5), 559–578.

Gourley, P. (2021). Curb appeal: How temporary weather patterns affect house prices. *The Annals of Regional Science*, *67*(1), 107–129.

Guren, A. M. (2018). House Price Momentum and Strategic Complementarity. *Journal of Political Economy*, *126*(3), 1172–1218.

Han, L., & Strange, W. C. (2015). The microstructure of housing markets: Search, bargaining, and brokerage. In G. Duranton, J. V. Henderson, & W. C. Strange (Eds.), *Handbook of regional and urban economics* (pp. 813–886, Vol. 5). Elsevier.

Han, L., & Strange, W. C. (2016). What is the role of the asking price for a house? *Journal of Urban Economics*, *93*, 115–130.

Haurin, D. R., Haurin, J. L., Nadauld, T., & Sanders, A. (2010). List prices, sale prices and marketing time: An application to us housing markets. *Real Estate Economics*, *38*(4), 659–685.

Hayunga, D., Pace, R. K., Zhu, S., & Calabrese, R. (2024). Differential measurement error in house price indices. *The Journal of Real Estate Finance and Economics*, 1–25.

Hirshleifer, D., & Shumway, T. (2003). Good day sunshine: Stock returns and the weather. *The journal of Finance*, *58*(3), 1009–1032.

Horowitz, J. L. (1986). Bidding models of housing markets. *Journal of Urban Economics*, *20*(2), 168–190.

Hsieh, C.-T., & Moretti, E. (2019). Housing constraints and spatial misallocation. *American Economic Journal: Macroeconomics*, *11*(2), 1–39.

Jacobsen, B., & Marquering, W. (2008). Is it the weather? *Journal of Banking & Finance*, *32*(4), 526–540.

Johnson, E. J., & Tversky, A. (1983). Affect, generalization, and the perception of risk. *Journal of personality and social psychology*, *45*(1), 20.

Kahneman, D., & Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, *47*(2), 263.

Kahneman, D., & Tversky, A. (1984). Choices, values, and frames. *American Psychologist*, *39*(4), 341–350.

Kamstra, M. J., Kramer, L. A., & Levi, M. D. (2003). Winter blues: A SAD stock market cycle. *American Economic Review*, *93*(1), 324–343.

Kamstra, M. J., Kramer, L. A., & Levi, M. D. (2015). Seasonal variation in treasury returns. *Critical Finance Review*, *4*(1), 45–115.

Kamstra, M. J., Kramer, L. A., Levi, M. D., & Wermers, R. (2017). Seasonal asset allocation: Evidence from mutual fund flows. *Journal of Financial and Quantitative Analysis*, *52*(1), 71–109.

Kelly, P. J., & Meschke, F. (2010). Sentiment and stock returns: The sad anomaly revisited. *Journal of Banking & Finance*, *34*(6), 1308–1326.

Kőszegi, B., & Rabin, M. (2006). A model of reference-dependent preferences. *The Quarterly Journal of Economics*, *121*(4), 1133–1165.

Lamorgese, A. R., & Pellegrino, D. (2022). Loss aversion in housing appraisal: Evidence from Italian homeowners. *Journal of Housing Economics*, *56*, 101826.

Lancaster, K. J. (1966). A new approach to consumer theory. *Journal of Political Economy*, *74*(2), 132–157.

Lantz, B. (2023). *Machine Learning with R* (4th ed.). Packt Publishing Ltd.

Leamer, E. E. (2007, September). *Housing is the business cycle* (Working Paper No. 13428). National Bureau of Economic Research.

Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. *Advances in neural information processing systems*, *30*.

Magnusson, A. (2000). An overview of epidemiological studies on seasonal affective disorder. *Acta Psychiatrica Scandinavica*, *101*(3), 176–184.

Magnusson, A., & Boivin, D. (2003). Seasonal affective disorder: An overview. *Chronobiology international*, *20*(2), 189–207.

Mersch, P. P. A., Middendorp, H. M., Bouhuys, A. L., Beersma, D. G., & van den Hoofdakker, R. H. (1999). Seasonal affective disorder and latitude: A review of the literature. *Journal of Affective Disorders*, *53*(1), 35–48.

Moen, E. R., Nenov, P. T., & Sniekers, F. (2021). Buying first or selling first in housing markets. *Journal of the European Economic Association*, *19*(1), 38–81.

Ngai, L. R., & Sheedy, K. D. (2020). The decision to move house and aggregate housing-market dynamics. *Journal of the European Economic Association*, *18*(5), 2487–2531.

Novy-Marx, R. (2009). Hot and cold markets. *Real Estate Economics*, *37*(1), 1–22.

Piazzesi, M., Schneider, M., & Stroebel, J. (2020). Segmented housing search. *American Economic Review*, *110*(3), 720–759.

Pissarides, C. A. (2000). *Equilibrium unemployment theory*. MIT press.

Rekkas, M., Wright, R., & Zhu, Y. (2022, July). *How Well Does Search Theory Explain Housing Prices?* (Working paper). Available at SSRN 3706329.

Rosen, S. (1974). Hedonic prices and implicit markets: Product differentiation in pure competition. *Journal of Political Economy*, *82*(1), 34–55.

Saunders, E. M. (1993). Stock prices and wall street weather. *The American Economic Review*, *83*(5), 1337–1345.

Schwarz, N. (2011). Feelings-as-information theory. In P. A. M. Van Lange, E. T. Higgins, & A. W. Kruglanski (Eds.), *Handbook of theories of social psychology* (pp. 289–308, Vol. 1). Sage.

Schwarz, N., & Clore, G. L. (1983). Mood, misattribution, and judgments of well-being: Informative and directive functions of affective states. *Journal of personality and social psychology*, *45*(3), 513.

Schwarz, N., & Clore, G. L. (2007). Feelings and phenomenal experiences. In A. W. Kruglanski & E. T. Higgins (Eds.), *Social psychology: Handbook of basic principles* (2nd ed., pp. 385–407). Guilford, New York, NY.

Statistics Norway. (2018). Boligtype- og standard for personer, etter familiefase (prosent) (avslutta serie) 2012 - 2018. Retrieved October 17, 2024, from https://www.ssb.no/statbank/table/09761

Statistics Norway. (2024a). Befolkning og endringer, etter statistikkvariabel og år. Retrieved October 11, 2024, from https://www.ssb.no/statbank/table/06913

Statistics Norway. (2024b). Boliger, etter region, statistikkvariabel og år. Retrieved October 11, 2024, from https://www.ssb.no/statbank/table/06265/

Statistics Norway. (2024c). Makroøkonomiske hovedstørrelser. Ujustert og sesongjustert 1978K1 - 2024K2. Retrieved October 11, 2024, from https://www.ssb.no/statbank/table/09190

Statistics Norway. (2024d). Omsetning av boligeiendommer med bygning i fritt salg, etter statistikkvariabel og år. Retrieved October 11, 2024, from https://www.ssb.no/statbank/table/06726/

Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive psychology*, *5*(2), 207–232.

Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, *5*, 297–323.

Wheaton, W. C. (1990). Vacancy, search, and prices in a housing market matching model. *Journal of political Economy*, *98*(6), 1270–1292.

Yavas, A., & Yang, S. (1995). The strategic role of listing price in marketing real estate: Theory and evidence. *Real estate economics*, *23*(3), 347–368.

Zhou, T., Clapp, J. M., & Lu-Andrews, R. (2022). Examining omitted variable bias in anchoring premium estimates: Evidence based on assessed value. *Real Estate Economics*, *50*(3), 789–828.

2 On Loss Aversion in Housing Markets: Evidence from Norway

On Loss Aversion in Housing Markets: Evidence from Norway*

Andreas Eidspjeld Eriksen†

November 11, 2024

Abstract

This paper investigates whether sellers in the housing market are loss averse, which is argued to be the source for the positive correlation between prices and volume in housing markets. Using a sample of repeated purchases-to-listings, and measuring sellers' price expectation with observed appraisal values, the findings suggest that sellers do not exhibit loss averse behavior in their list price setting. In contrast, results from the standard approach of measuring expectations with hedonic predictions of selling prices show significant estimates of the effect. The contribution is two-fold. First, the reason for this difference in estimates is found to be due to omitted variable bias from unobserved time-invariant and time-varying heterogeneity in prices. This highlights the issue of omitted variable bias that arises when using substitutes that introduces differential measurement error. Second, institutional factors may influence whether sellers act in accordance with being loss averse.

Keywords: *Loss aversion; Housing market; Price expectation; Measurement error*
JEL classification: *D91; R21; R31*

1 Introduction

The phenomenon of a positive price-volume correlation during housing market busts has been previously been explained by sellers not willing to sell their housing units at the current expected

*The paper was presented at the 44th Annual Meeting of the Norwegian Association of Economists (Forskermøtet 2022), the 98th Annual Conference of the Western Economic Association International, and the 2023 Housing and Urban Research Workshop at Oslo Metropolitan University. I am grateful to André Kallåk Anundsen, Anne Wenche Emblem, Cloé Garnache, Svenn Jensen, Erling Røed Larsen, Plamen Nenov, Dag Einar Sommervoll, and seminar participants at Norges Bank for helpful comments. I also want to thank Eiendomsverdi AS for providing the transaction data to this study.

†School of Economics and Business, Norwegian University of Life Sciences and Housing Lab, Oslo Metropolitan University. E-mail: andrease@oslomet.no

market prices (Genesove & Mayer, 2001). This leads to longer time-on-market and list prices that are larger than the expected market prices, which accumulates to a larger unsold inventory. Conversely, during booms, housing markets are characterized by shorter time-on-market and smaller unsold inventories. A possible explanation for this kind of seller behavior is that they are loss averse. The sellers expect losses in utility from selling their housing units when their market price expectations are below their reference points. When facing prospective losses, sellers should put list prices above their price expectations, hoping to match with buyers willing to pay the higher prices.[1]

This paper asks whether sellers in the Oslo housing market are nominally loss averse in their list price choices. To investigate sellers' loss aversion, I use data from the Oslo housing market and propose a new measure of sellers' own market price expectations, namely the appraisal values provided by surveyors shortly before the units are listed on the market. In short, I find that sellers does not exhibit loss averse behavior in their list price setting, but inference depends on the institutional context in which these sellers operate. I also find that the standard way of measuring the price expectations of sellers can lead to biased results beyond what has been documented in previous studies. Meaning, the possibility of inference relies on the available information about price expectations.

Loss aversion is rooted in prospect theory introduced by Kahneman and Tversky (1979). They present an alternative model of expected utility which captures behavioral effects when economic agents face risk. Loss averse people dislike losing more than they like gaining, so that their decision making is affected by their aversion against losing. In the housing market literature, the seminal paper using micro-level data of listings and transactions to identify the effect is by Genesove and Mayer (2001) (hereafter GM). They present an empirical framework to identify sellers' loss aversion, which consists of modeling the list price with two explanatory variables: the expected price and the prospective loss. Thus, the loss aversion effect in the market is estimated on the intensive margin, specifically in the choice of list price. GM find a significant loss aversion effect in the 1990s Boston condominium market. Using a sample of repeated purchases-to-listings in the Oslo housing market in the 2005 to 2020 period, I estimate the GM model with the extension of adding the prospective gain to the model, as proposed by Bokhari and Geltner (2011). The GM approach is to use hedonic predictions as substitutes for the expected market price, resulting in a significant loss aversion effect when replicated using the Oslo sample. The estimates suggest that a seller facing a prospective loss of 10 percent is associated with increasing the list price by between 4.2 and 6.1 percent, ceteris paribus. In comparison, in the Boston condominium market, GM estimates the effect to be between 2.5 and 3.5 percent.

[1]This is a fishing strategy: higher list prices are put at the expense of longer time-on-market. See Figure D.1 in Appendix D for a visualization of the relationship between prices and time-on-market in Oslo.

In studies of loss aversion in housing markets, hedonic predictions or indices are used to capture the unobserved expectation about market price. Predictions do not account for all of the variation in the true values. When measuring expectations of prices using predictions of actual selling prices, this introduces measurement error in the form of prediction error, while the predicted variable itself is differing from the expectations. Therefore, the measurement error consists of three unobserved parts, the time-invariant heterogeneity, the time-varying heterogeneity, and the deviation between the expected market price and the actual selling price. These are all omitted variables that may lead to bias in the estimates of the effect, which is apparent from the two point estimates argued to be the lower and upper bound of the true effect. A recent paper investigating loss aversion and addressing the unobserved heterogeneity problem is Andersen et al. (2022). They use a structural model approach to quantify the effect of loss aversion among Danish homeowners, and assess the issues with measurement error by implementing different approaches of predicting prices, giving a kink at zero prospective gains for all approaches. Yet, some sources of bias still remains even with their implementations. Their proposed alternative methods reveal the difficulties of inference using prediction-based models, thus motivating the use of an observed price rather than a predicted price.[2]

In this paper, I propose capturing price expectations by using appraisal values. Until mid-2016, sellers and realtors used appraisal values as tools to determine list prices and as supplements in marketing to indicate potential market values. Appraisal values were supplied by surveyors when evaluating the technical condition of the units.[3] There are several advantages of using the appraisal values compared to hedonic predictions. The appraisal values are decided upon by experts who visit the housing units, thus, when considering appraisal values as predictions, more of the heterogeneity absent from any register data are available to the surveyors, such as general quality and view. Moreover, the fact that sellers are supplied with appraisal values is a key institutional factor that plays into the list price setting for two reasons. First, the sellers' expectations may shift when given expert opinions (Northcraft & Neale, 1987; Tversky & Kahneman, 1974). Second, knowing that buyers observe the appraisal values, sellers feel constrained in their list price setting, so that they do not put list prices above their respective appraisal values. Therefore, the low degree of asymmetric information, and that sellers update their expectations according to expert opinions, may mitigate the effect from potential loss aversion on list prices. This is evident from the results, which suggest that there is no such loss aversion effect on list prices once the appraisal values are taken into account.

[2]For instance, as presented in Andersen et al. (2022), autocorrelated unobserved heterogeneity could still bias the estimates when applying the repeated sales models. Also, in the case of only a single repeated sale, the feasible variables of interest reduces to be the same as studied by Bracke and Tenreyro (2021), meaning prospective gains/losses in aggregate prices. Hence, the choice of prediction method is important for inference of the causal relationship.

[3]The practice of using surveyors is very common in Norway. In mid-2016 surveyors stopped providing appraisal values but surveyors are still hired for technical evaluations.

In turn, a new question arises: why is there a difference between the results from estimating the hedonic-based and appraisal-based models? In principle, the appraisal-based model should yield an estimate between the lower and upper bound from estimating the hedonic-based model. A possible reason for the gap between estimates is the substitution of predictions as explanatory variables. If appraisal values are preferred over predictions, and the preferred estimation approach suggests that sellers are not affected by loss aversion in their list price setting, then concerns about biases in estimates should be greater when estimating the hedonic-based model. To provide some insights into the difference between the two, I build upon the reduced form investigation of GM, who assume that unobserved heterogeneity is time-invariant. In theory, relaxing this assumption, by allowing for unobserved time-varying heterogeneity, could explain the results. Therefore, my approach to identifying why the estimates differ is by step-wise dealing with both time-invariant and time-varying unobserved heterogeneity, primarily in the model that substitutes expectations with predictions.

Starting with estimating models with unit fixed effects the estimated loss aversion effect becomes smaller. However, this approach does not account for time-varying unobserved heterogeneity. Estimating a model with hedonic predictions that take more information into account results in smaller estimates of the effect. Estimation with both improved predictions and unit fixed effects yields an even smaller estimate: a seller facing a prospective loss of 10 percent is associated with increasing the list price by between 1.9 and 2.6 percent, ceteris paribus. These results indicate that unobserved heterogeneity indeed is important for inference, but none of the methods fully address all the unobserved heterogeneity. Meaning, unobserved time-varying heterogeneity is present and biases the estimates.

Moreover, an additional source of measurement error is that it is not obvious that the substitutes for price expectations should be predictions of selling prices, because the price expectations cannot be outcomes of unrealized selling prices at the time of listing. I obtain fitted values of appraisal values to test whether they give different results as substitutes than predictions of selling prices. The estimates suggest that the choice between selling prices and appraisal values cannot alone be what creates the difference between the loss aversion estimates.

Although the results show that omitted variable bias is present, it is not obvious through which channels the biases enter. The right-hand side of the model has three key variables. First, the prospective loss and gain, which are the variables of interest. Second, a control variable for the price expectation of the seller, referred to as the *main list price determinant*. While this is included as a control variable, it is also used in constructing the prospective loss and gain variables. Therefore, all of the key right-hand side variables suffer from measurement error when using substitutes for the true expectations. This motivates a step-wise investigation to trace which part of the model that is more sensitive to the two alternative substitutes in question. Models with different substitutes for the main list price determinant, as well as for constructing the prospective loss and gain, are estimated. For instance, a model uses hedonic predictions as the main list price determinant and appraisal values as substitutes in the prospective loss and gain. These models are estimated to determine whether measurement error in the prospective terms, the main list price determinant, or both, is producing the difference between loss aversion estimates. The estimates suggest that there is a spillover effect from the main determinant to the prospective terms. Moreover, when adding more random noise to this main determinant while keeping the prospective terms unchanged, the hedonic-based model yields larger estimates. Together with evidence from utilizing unit fixed effects and providing more information in the hedonic model, I find that time-varying unobserved heterogeneity is an important component of the measurement error in predictions, which ultimately biases the estimates.

The measurement error introduces differential (Berkson) error (Berkson, 1950; Haber et al., 2021) in the loss aversion model, resulting in omitted variable bias. Differential measurement error arises when substituting an explanatory variable with a proxy that has measurement error that is correlated with other explanatory variables while also being important for the dependent variable. In total, the direction of the bias is close to being untractable when the measurement error consists of the time-varying unobserved heterogeneity and all the important explanatory variables suffer from this measurement error.[4]

The findings nuance the current literature on the subject matter. By considering unobserved heterogeneity in prices that vary over time, it is evident that estimates of loss aversion among housing market sellers may be upward biased to a greater extent than previously believed. This is relevant for the existing literature on loss aversion in housing markets (e.g., Andersen et al., 2022; Anenberg, 2011; Bokhari and Geltner, 2011; Einiö et al., 2008; Genesove and Mayer, 2001; Lamorgese and Pellegrino, 2022). Specifically, the findings suggest that sellers do not exhibit loss aversion in their list price setting in the Oslo housing market. In this housing market, sellers know that buyers will observe the appraisal values, and these appraisal values should have influenced their expectations about selling prices. Therefore, sellers could have felt constrained in their list price decision, and this may have overpowered their urge to act according to impulses caused by them being loss averse. Therefore, whether loss averse behavior is prevalent among sellers may depend on institutional factors.

The rest of the paper is structured as follows. Section 2 describes the data, the cleaning of the data, and the institutional background. Section 3 presents the identification framework and the estimation results from the Oslo market. This is followed by an investigation of the differing estimation results between the hedonic-based model and appraisal based model in Section 4. Robustness and sensitivity tests are presented in Section 5, and the conclusion in Section 6.

[4]The model has three crucial right-hand side variables, all of which suffer from the measurement error.

2 Data and Institutional Details

2.1 Description

The data used in this paper come from two sources. Housing transactions are provided by Eiendomsverdi AS, a private firm that provides market price predictions of real estate properties in commercial banks' mortgage portfolios and offers a tool for assessing values of other properties often used by financial advisors and realtors, among other services.[5] The dataset includes unique consistent identifiers for housing units and individual sellers and buyers, and covers the period 2003 to 2021 for the 20 largest municipalities in Norway, as measured by population size. The data include dwellings sold on the market, with no withdrawn listings. Other variables include the roles of each individual in each transaction and their ownership share, selling and list prices, dates for when the units are listed and sold, attributes, zip codes, dwelling type, and ownership type. Information about sellers and buyers comes from administrative data provided by Statistics Norway. The data cover the period 2005 to 2019 and have a panel structure. The variables included are different income, debt, and wealth variables, all being measured at the end of the year.

To avoid errors in the data to spill into estimation results, the data were cleaned. There are two main ownership types in Norway: co-ops and non-co-ops. In practice, the two do not differ much.[6] Typically, in larger cities, a larger share of apartments have a co-op ownership structure. Both of these types are included in the data. Before 2007, there were no organized registration of ownership of co-ops, but after 2007 all co-op units were registered in the Land Registry just like non-co-op units. Transactions of co-ops before 2007 are therefore removed from the data.

The data cleaning process continued as summarized in Table D.1 in Appendix D. The sample containing dwellings sold once, or more, are used to fit hedonic regressions in order to predict prices.[7] In order to get a repeated structure, the sample is reduced to contain only repeat purchases-to-listings, i.e., only listings in which the seller(s) are the same as the buyer(s) in the preceding sale of that same housing unit. Naturally, there are few observations in the first couple of years. Also, there is no debt data available for 2020, which is information needed for the analysis, therefore, the sample period is restricted to Q1 2005 through Q4 2020.

Previous studies typically use loan-to-value (LTV) to control for the effect of being equity-

[5]Eiendomsverdi AS is owned by some of the largest commercial banks in Norway. The value assessment tool is based on an automatic valuation model of house prices. Eiendomsverdi AS gathers information from different public sources and realtor firms, and combines these data to a unique commercial data set.

[6]The main difference between the two is the ownership structure. Owners of non-co-ops own their own housing units directly. In co-ops, the housing cooperative owns the building. Owners of co-op units hold shares in the housing cooperative, giving them the right to reside in their respective housing units. Also, buyers of co-op units does not pay stamp duty, while non-co-op units have a stamp duty tax of 2.5 percent.

[7]More information of these regressions are given in Appendix A.1.

constrained. The actual mortgages related to the dwellings are not available in the data, therefore, I use debt-to-value (DTV), defined as the total debt relative to an estimated market value.[8] The total debt consists of all registered debt aggregated on the sellers in each transaction, and it is measured at the end of the year preceding the listing. This introduces some lag to the debt but ensures a measure of the equity-debt situation at the time of listing.

The data include appraisal values, meaning market price estimates given by surveyors that have inspected the housing units prior to the listings. There is a structural break in these values in mid-2016, which is due to the transition among realtors from using appraisal values to predictions from an Automated Valuation Model. The few observations with appraisal values after June 2016 are removed from the sample when estimating models that use appraisal values.

This results in two samples: a sample used when estimating models that use hedonic predictions, and a sample used when estimating models that use appraisal values. The hedonic sample consists of 32,044 observations and covers the period January 2005 to December 2020, and the appraisal value sample consists of 16,111 observations and covers the period January 2005 to June 2016. Summary statistics of the repeated purchases-to-listings samples are presented in Table 1.[9] Note that all prices include common debt, if any, which is typically a shared debt among apartments in the same building.[10]

The model that is estimated is a list price model. Only the latest list prices are included in the data. This could affect the results since initial list prices may be higher in some cases.[11] Potentially, the may dampen the estimated effect because loss averse sellers should put higher list prices than others. Table 2 shows the distribution of repeated sales in the repeated sales data. The fact that many units are sold multiple times, such that they have multiple repeated sales, is a feature utilized for unit fixed effects, used to deal with unobserved time-invariant heterogeneity in the estimations.

[8]DTV is truncated from below at 0.85, so that $DTV^{trunc} = (DTV - log(0.85) - 1)^{+}$, with DTV being the difference between the log of the total debt and log of the market value. This difference is subtracted with $log(0.85) + 1$ (≈ 0.85), the log-equivalent to a 85 percent DTV cutoff. There is a current equity requirement of 15 percent, meaning a maximum LTV of 85 percent. The 0.85 cut-off on the DTV is chosen to capture those being close to this limit, but the DTV measures the debt-to-value situation relatively shortly before the listing. When hedonic predictions of log selling price are used in estimations, the log market value is substituted with the predicted log selling price. When appraisal values are used in estimations, the substitute is the log appraisal value.

[9]See Figure D.2 in Appendix D for histograms of annual observations for both samples, and Table D.2 for summary statistics of the non-repeated data and the raw repeated data.

[10]Common debt is usually a result of renovations of shared features of a building, such as the exterior. It is more common for co-ops.

[11]Based on data from the second largest Norwegian realtor company, DNB Eiendom, which include list price revisions, about 76 percent of transactions does not have revisions, while about 16 percent have one revision. These revisions could be registered before the actual listing on the market.

Table 1: Summary statistics

Variable	1st Qu.	Median	Mean	3rd Qu.
(A) *Hedonic sample (N=32,044, Jan 2005–Dec 2020)*				
List price (MNOK)	2.58	3.37	3.73	4.40
Selling price (MNOK)	2.74	3.50	3.88	4.55
Predicted price at listing (MNOK)	2.69	3.51	3.82	4.53
List-Predicted spread (%)	-9.83	-3.37	-2.71	3.76
Size (m^2)	50	63	68	79
TOM (days)	9	10	19	16
Holding time (weeks)	129	198	222	289
Apartment (%)			92.60	
Non-co-op (%)			57.25	
(B) *Appraisal sample (N=16,111 Jan 2005–Jun 2016)*				
List price (MNOK)	2.10	2.69	3.07	3.63
Selling price (MNOK)	2.27	2.85	3.24	3.82
Appraisal value (MNOK)	2.15	2.72	3.12	3.70
List-Appraisal spread (%)	-3.06	-0.38	-1.78	0.00
List-Predicted spread (%)	-10.77	-4.26	-3.34	3.04
Size (m^2)	50	64	70	82
TOM (days)	9	11	18	14
Holding time (weeks)	117	180	198	258
Apartment (%)			91	
Non-co-op (%)			66	

Notes: The table reports summary statistics for the two repeated purchases-to-listings samples from the Oslo housing market. Prices are reported in millions of Norwegian kroner, and include common debt. The reported predicted price are the exponential of the ones used in the analysis, meaning they should be smaller than if the hedonic specification was in a linear form due to Jensen's inequality for concave functions. Holding time is the number of weeks between the previous date of sale and the date of listing. TOM denotes the time-on-market measured as the number of days between the listing and the sale.

2.2 Market characterization and institutional background

The Norwegian housing market can be characterized as transparent. The market is regulated in order for buyers to have as much information as possible about housing units for sale. When sellers hire realtors,[12] the norm is that the realtors hire surveyors. Surveyors visit the units and provide technical reports that describe their condition. This is done before listing the units on the market, and these reports are included as attachments in the sales-prospectus provided to potential buyers. The idea is to hire a professional surveyor to reduce the risk burden of selling a unit without being fully transparent of potential hard-to-observe qualities, either good or bad. These surveyors are the same that supplied appraisal values until mid-2016.

[12]The Norwegian Real Estate Association reports that 98% of on-market sales use realtors.

Table 2: Frequencies of repeated sales

Number of repeated sales	1	2	3	4	5	6
Hedonic subsample	15,088	5,393	1,576	308	36	5
Appraisal value subsample	10,299	2,312	344	34	4	

Notes: The table presents total number of repeated sales of unique units in the two samples. For instance, if a unit is observed sold three times, and the sellers and buyers match between observations, then this unit has two repeated sales.

The median (mean) TOM in days is only 10 (19) days in the larger hedonic sample, suggesting two interesting aspects of the market. First, the TOM is very low compared to other countries. For instance, median TOM in the U.S. housing market is reported to be 69 and 66 days for 2019 and 2020, respectively (Realtor.com, 2022). It is common to conduct public showings about a week after the initial listing, hence, the median suggests that it is relatively common to successfully sell in the days after these showings. Second, some units are harder to sell. Even after dealing with outliers, the distribution is right-skewed.

In short, a usual sales process in the Norwegian housing market goes as follows. A realtor is hired by a seller. Potential buyers do not hire realtors. The realtor prepares all paperwork and marketing of the unit, including hiring a surveyor and a photographer. After receiving all required documentation, including the report from the surveyor about the technical condition and a market price estimate (the appraisal value), the unit is listed with a list price and a detailed account of the unit. The internet marketplace Finn.no is the most commonly used place for listing real estate. Public showings are conducted, and the price is determined through a bidding process that resembles an open ascending bid auction.

Adding to this account of the sales process, one may question whether it is the seller or the realtor that decides the list price.[13] When a seller hires a realtor, the seller makes a decision based on a presentation of a total package given by the realtor. The package includes an estimate of the total cost of the realtor's service and a rough estimate of the value of the unit. Total cost includes the commission based on a suggested commission rate. The rough estimate can either become or influence the price expectation of the seller, therefore, it may be affect the choice of list price. However, if the seller is getting in contact with realtors in the first place, the seller probably already has an idea of what the unit is possibly worth. For markets were appraisal values were used, such as the Oslo market, the value estimate presented by the realtor was meant to inform the seller about the commission cost. The surveyor would later provide the seller and realtor with the appraisal value, which should be a stronger determinant for the list price. One may think of this as the seller updating her expectation multiple times,

[13]In the Norwegian institutional setting, Anundsen et al. (2022) model the list price decision as the outcome of Nash bargaining between sellers and realtors.

in vein of Bayesian updating or anchoring-and-adjustment. Even if the seller had some initial expectation, this may shift based on the newly provided information. The appraisal value is the latest price estimate presented to the seller, potentially being the final source of shifting the expectation before listing.

3 Identification framework

3.1 Genesove and Mayer (2001) framework

The models applied in this paper are based on the model presented by Genesove and Mayer (2001). Their identification framework captures prospective losses from two sources: aggregate prices and a micro-level, purchase-specific term that reflects whether the sellers paid too much or too little for the units they are now selling. By adopting the GM notation their ideal model takes the following form:

$$\begin{aligned} L_{ist} &= \alpha_0 + \alpha_1 \mu_{it} + m_l LOSS^*_{ist} + \epsilon_{it} \\ &= \alpha_0 + \alpha_1 \mu_{it} + m_l \left(P_{is} - \mu_{it}\right)^+ + \epsilon_{it}, \end{aligned} \tag{1}$$

in which L_{ist} denotes the log list price of unit i at the time of listing t. This list price is explained by the log expected market price, μ_{it}, and the nominal prospective loss that the seller faces when listing the unit, $LOSS^*_{ist}$. The prospective loss is defined as the difference between the log selling (purchasing) price when the unit was sold last time, P_{is}, and the log expected market price at the time of listing, truncated from below at zero. Thus, the $LOSS$-variable only takes non-negative values, being positive for prospective losses. When μ_{it} enters the right-hand side as a control variable, this is referred to as the *main list price determinant*. Using this specification implies the assumption that sellers' reference points are what they paid for their housing units, so that they evaluate losses and gains in nominal terms from this reference point.

$$L_{ist} = \alpha_0 + \alpha_1 \left(X_i \beta + \delta_t + v_i\right) + m_l \left(\delta_s - \delta_t + w_{is}\right)^+ + \epsilon_{it}, \tag{2}$$

Further, equation (2) gives the reduced form of the ideal model. Here, the log expected market price takes the linear reduced form $\mu_{it} = X_i \beta + \delta_t + v_i$, in which X_i is a vector of the observed time-invariant attributes, while δ_t is the quarter of listing, and v_i is unobserved time-invariant heterogeneity. The log selling price when the seller purchased the unit can be reduced in the same manner, but adding a term for whether the seller paid above, or below, the expected market price, so that $P_{is} = \mu_{is} + w_{is}$, in which $w_{is} > 0$ means that the seller paid above the log expected market price. Finally, ϵ_{it} is a disturbance which is assumed to have a conditional mean of zero.

To estimate the relationship, and because the true expected price is unobserved, μ_{it} can be substituted with a prediction of the log selling price. Predictions from a hedonic model only accounts for observed variables: $\hat{P}_{it} = X_i\hat{\beta} + \hat{\delta}_t$. This means that both explanatory variables in equation (1) are substituted with proxies. GM discuss implications of substituting expectations with predictions, concluding that the resulting estimate $\hat{m}_l$ could be upward biased. Adding the residual from the selling price prediction of the previous sale, $\epsilon_{is} = v_i + w_{is}$, is argued to result in a downward biased $\hat{m}_l$. Some studies (e.g., Beggs and Graddy, 2009; Bokhari and Geltner, 2011; Graddy et al., 2022; Lamorgese and Pellegrino, 2022) utilize this lagged residual as a quality control variable to account for unobserved heterogeneity. This is reasonable when the residuals across multiple sales do not show tendencies of reversion and when the selling prices are likely not influenced by over- or underpayment (w_{is}) relative to the expected market price. Furthermore, unobserved attributes and any omitted functional forms of observed and unobserved attributes are assumed to be constant over time, thus being part of v_i. Nonetheless, it seems reasonable to expect the true coefficient to be somewhere in the range between the two boundary values, $m_l \in (\hat{m}_l^{LOWER}, \hat{m}_l^{UPPER})$, when estimating the relationship with and without lagged hedonic residual.

3.2 Bokhari and Geltner (2011) extension

Bokhari and Geltner (2011) (hereafter BG) expand the GM framework by including the prospective gain as an additional explanatory variable of interest, defined as $GAIN^*_{ist} = (P_{is} - \mu_{it})^-$.[14] There are mainly two reasons for including the gain. First, in prospect theory, loss aversion is the kink in the value function at the reference point (Kahneman & Tversky, 1984, p. 342). By only including prospective losses in the model, this assumes that the coefficient on prospective gain (m_g) is zero. If this was the case, the estimate of m_l could identify the effect. Sellers facing gains may still put lower list prices compared to situations where they face neither losses nor gains.[15] Hence, the difference in coefficients, $m_l - m_g$, is what identifies loss aversion among sellers.

Second, if the prospective gain is important for the list price, it is an omitted variable contributing with positive bias on the loss coefficient: prospective loss and gain correlates positively with each other.[16] A low share of prospective losses compared to gains makes the

[14] Omitting this variable is done deliberately by GM: the sum of the loss and gain nests the residual proxy that is used to downward bias the coefficient estimate of m_l (Genesove & Mayer, 2001, p. 1241). Meaning, $LOSS_{ist} + GAIN_{ist} = \delta_s - \delta_t + v_i + w_{is}$, and the residual control is $v_i + w_{is}$. This is a reasonable concern in light of multicollinearity. However, it would be a greater concern if the loss and gain were added together as a compound of the two, and not included separately as suggested. By separating the two, this probably will result in different coefficients on the two, and the nested control will enter with a discontinuity. BG also argue that adding prospective gain together with the residual, ϵ_{is}, solves the GM identification problem.

[15] There can be many reasons for this, such as those facing gains being willing to trade off some gains for a faster sale.

[16] Let X be a random variable, then $X := (X)^+ + (X)^- = X_1 + X_2$. Because $X_1X_2 = 0$, hence $E(X_1X_2) = 0$,

bias from omitting prospective gains larger, a feature that should be present in rising housing markets.

The feasible model to be estimated that accounts for the prospective gain, without the lagged residual as a control variable, is presented in equation (3) and its corresponding residual in (4). A review of potential biases when estimating this specification is presented in Appendix B.1.[17]

$$L_{ist} = \alpha_0 + \alpha_1 \hat{P}_{it} + m_l LOSS_{ist} + m_g GAIN_{ist} + \eta_{it} \tag{3}$$

$$\eta_{it} = \alpha_1(\mu_{is} - \hat{P}_{it}) + m_l \left(LOSS^*_{ist} - LOSS_{ist}\right) + m_g \left(GAIN^*_{ist} - GAIN_{ist}\right) + \epsilon_{it} \tag{4}$$

By adding the prospective gain, it is possible to consider whether there is no behavioral effect ($m_l = m_g = 0$), a pure reference-dependence effect ($m_l = m_g \neq 0$), or a loss aversion effect with, or without, reference-dependence ($m_l > m_g \geq 0$). For the loss aversion effect, the kink at the reference point should manifest in a stronger effect of facing a prospective loss: $m_l > m_g$. The intuition in this framework is not as straight forward as it may seem. Here, loss aversion is identified if loss averse sellers facing prospective losses actually put higher list prices than they otherwise would. This hinges on the assumption that sellers (unknowingly) believe that putting higher list prices could lead to higher selling prices. Applying the framework means that the "otherwise"-case does not only imply zero prospective loss but potential prospective gain. Using price predictions to measure prospective losses and gains should, without rounding, not result in any cases where these are zero. If sellers are not loss averse, but purchasing prices work as anchors, they would care about deviations from this reference point. Then they should put higher list prices when facing prospective losses and lower list prices when facing prospective gains, with the effect being at the same magnitude for losses and gains.

To get an idea of when prospective losses occur, it can be seen from Figure 1 how they appear in periods after small price peaks. Interestingly, the losses seem to appear at the highest frequency when the index recover to the previous peak level. Although the main driving force behind the occurrence of prospective losses should be reductions in the aggregate price level, several cases with over-payments ($w > 0$) in combination with small changes in aggregate prices may be the reason for the post back-to-peak cases.

---

the asymptotic covariance is $Cov(X_1, X_2) = -E(X_1)E(X_2)$. Prospective gain has negative expectation so that the loss-gain covariance is positive. And from what is commonly known about omitted variable bias, the smaller the variance of the included variable, the stronger the bias.

[17] The specification is actually the same as, e.g., the right-hand-side in Lamorgese and Pellegrino (2022). They estimate $SB_{it} = \alpha_0 + \alpha_1 \hat{P}_{it} + m_l LOSS_{ist} + m_R R_{ist} + \eta_{it}$ with SB_{it} being the homeowner's stated belief about market price and $R_{ist} = LOSS_{ist} + GAIN_{ist}$. Therefore, the coefficient on prospective loss is the same as $m_l - m_g$ in equation (3). Graddy et al. (2022) also use this specification when estimating loss aversion and reference dependence in the art market.

Figure 1: Price index and number of prospective losses

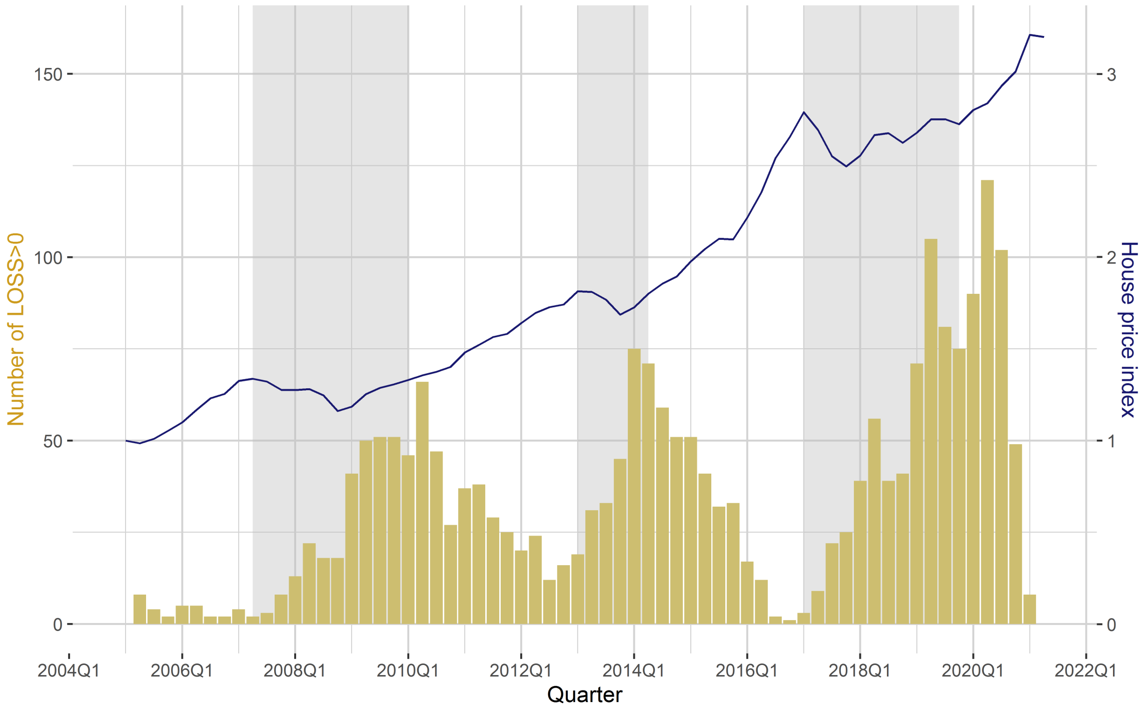

Notes: The figure presents a price index based on a hedonic prediction of house prices using year-quarter of sale dummies (not the quarter of listing dummies). The prospective losses are found using a backward-looking hedonic regression approach. The shaded areas begin at peaks and end when the index returns to the peak-level.

3.3 Substituting with appraisal values

A recent contribution to the literature on how to partly deal with the problem of unobserved heterogeneity comes from Clapp and Zhou (2020) and Zhou et al. (2022). They show that substituting the expected market price with their normalized assessed value (hereafter NAV) yields lower estimates of m_l without using the residual control. A potential problem with using tax-related assessed values is that these assessments are conducted on a 5-year cycle. With a fast moving market these suffer from being lagged, potentially making them less representable for the true expectations. However, the benefit is that the assessors observe more than what the econometrician observes, so that variation in v_i is reduced. The normalization is done by adding town-year average log selling prices with the deviation between individual and town-year average of log assessed value. By doing this normalization, the NAV for housing units in each town-year differs only due to the assessed value deviation, which is argued to deal with some of the issues with the lagged nature of assessed values.

In the Norwegian context, the NAV procedure on appraisal values is not necessary. The surveyors visit the units, and the market values are appraised a few weeks before the listing.

Because surveyors observe more than the econometrician, the appraised values are suitable substitutes for expected market prices. They are also likely to influence or even anchor sellers' expectations, thereby addressing a different aspect of the issue related to measuring true expectations with predicted selling prices: selling prices are not the same as expectations about market prices.[18]

If not invoking the assumption of constant unobserved heterogeneity, the ideal specification becomes:

$$\begin{aligned} L_{ist} &= \alpha_0 + \alpha_1 \mu_{it} + m_l LOSS^*_{ist} + m_g GAIN^*_{ist} + \epsilon_{it} \\ &= \alpha_0 + \alpha_1 \left(X_i \beta + \delta_t + v_{it} \right) + m_l \left(\delta_s - \delta_t + v_{is} - v_{it} + w_{is} \right)^+ \\ &\qquad + m_g \left(\delta_s - \delta_t + v_{is} - v_{it} + w_{is} \right)^- + \epsilon_{it}, \end{aligned} \tag{5}$$

so that the prospective loss and gain depend on the change in aggregate prices, the change in unobserved heterogeneity, and whether the seller paid above the expected market price. Denote the log appraised value at the time of listing as $P^{AV}_{it} = X_i\beta + \delta_t + \tilde{v}_{it}$, with $\tilde{v}_{it}$ capturing the additional heterogeneity observed by the surveyor. The feasible specification and its residual are as follows, with the notation kept unchanged for simplicity:

$$\begin{aligned} L_{ist} &= \alpha_0 + \alpha_1 P^{AV}_{it} + m_l LOSS^{AV}_{ist} + m_g GAIN^{AV}_{ist} + \eta_{it} \\ &= \alpha_0 + \alpha_1 (X_i \beta + \delta_t + \tilde{v}_{it}) + m_l \left(\delta_s - \delta_t + v_{is} - \tilde{v}_{it} + w_{is} \right)^+ \\ &\qquad + m_g \left(\delta_s - \delta_t + v_{is} - \tilde{v}_{it} + w_{is} \right)^- + \eta_{it} \\ \eta_{it} &= \epsilon_{it} + \alpha_1 \psi_{it} + m_l \left(LOSS^*_{ist} - LOSS^{AV}_{ist} \right) + m_g \left(GAIN^*_{ist} - GAIN^{AV}_{ist} \right). \end{aligned} \tag{6}$$

$LOSS^{AV}_{ist}$ and $GAIN^{AV}_{ist}$ represent the prospective loss and gain, respectively, measured by substituting μ_{it} with P^{AV}_{it}. ψ_{it} is the difference between the actual expected market price and the appraisal value, making it the mistake made by the surveyor. The additional information observed by the surveyor ($\tilde{v}_{it}$) is part of v_{it}.[19] This means that the unobserved part is smaller for the surveyor than for the econometrician, so that $E|\psi_{it}| \leq E|v_{it}|$. Also, it should be reasonable that surveyors put different weights on the attributes observed by the econometrician and the aggregate price level, δ_t. However, deviations in functional form and implicit pricing are nested in $\tilde{v}_{it}$, so that the X_i and δ_t notation remains. The surveyor mistake, ψ_{it}, does not enter $LOSS^{AV}_{ist}$ or $GAIN^{AV}_{ist}$. Yet, by splitting up the unobserved heterogeneity at the previous sale, $v_{is} = \tilde{v}_{is} + \psi_{is}$, it is clear that autocorrelation in surveyor errors could bias estimates of m_l and m_g upwards. The difference between ideal and feasible prospective loss and gain should now be smaller in magnitude than for the hedonic equivalent, because more of what was unobserved

[18] Whether appraisal values should be preferred over hedonic selling price predictions in the list price model is discussed in Appendix C.1.

[19] ψ_{it} comes from the log difference of the prices: $\mu_{it} - P^{AV}_{it} = v_{it} - \tilde{v}_{it} = \psi_{it}$.

is now observed. The surveyor, who visits the house, and risks legal consequences if doing a poor job, should observe more of a unit's characteristics. Therefore, a reasonable assumption may be that ψ_{it} has a conditional mean of zero. An even stronger assumption is to impose that $\psi_{it} = 0$, making the residual η_{it} to collapse to ϵ_{it}, and ultimately making the ideal specification to be identified.[20]

Finally, substituting with the appraisal value highlights another implied assumption when estimating equation (3). If the expectation about market value is shared with prospective buyers, then they would probably be less willing to pay a list price that exceeds the expectation. Importantly, this depends on whether the price expectation is common knowledge, or at least partly so. This implies that whether the true coefficient m_l in equation (5) is positive, and different from m_g, requires that μ_{it} is available exclusively to the seller. Otherwise, the seller will know that buyers are observing a list price above the market price expectation, potentially making the seller less prone to increase the list price. This may not be the case for m_g: if the seller expects to gain from a sale, she might set a lower list price to reduce the time-on-market, because she is more willing to accept such a trade-off.

3.4 Results

The results of estimating the feasible model in equation (1) are provided in the first two columns of Table 3. The estimate of the upper bound is 1.49 and the lower bound is 0.43. The interpretation of the former is as follows: a seller facing a 10 percent increase in prospective loss is associated with increasing the list price by 14.9 percent, all else equal. These high estimates can be explained by the relatively few observed prospective losses in the sample. If there are omitted variables causing bias, this may be amplified by the low variation in prospective losses in the Oslo sample.[21] Moreover, people who face prospective losses while knowing this makes them outliers, compared to current or historical majority, may feel stronger about these losses than if such losses were more common. This argument assumes that loss aversion is amplified by paying too much or by being a losing outlier, which should be captured by higher values of w_{is}.

The two last columns in Table 3 are results from estimating the model with the BG extension. Compared to first two columns, estimation of the relationship in equation (3) is done without separating the predictions into a base price and an index,[22] and using price predictions

[20] The distribution of prospective appraisal-based losses are presented in Figure D.3.

[21] These are results of replicating the first two columns of Table 2 in Genesove and Mayer (2001) using the Oslo repeated purchase-to-listing data, and controlling for DTV rather than LTV. For the 1990s Boston condominium market, GM found that the effect is between 0.25 and 0.35, while replication in the Danish housing market by Andersen et al. (2022) resulted in an effect between 0.47 and 0.56. Although not stated explicitly, both these markets should have a higher share of positive prospective losses than the Oslo sample. It can also be argued that there are cultural differences that can partly explain the large differing results.

[22] The GM replication without this separation yielded very similar results.

Table 3: Replicating GM and BG extension

	GM		BG	
	(1)	(2)	(3)	(4)
LOSS	1.490***	0.435***	0.924***	0.489***
	(0.102)	(0.078)	(0.075)	(0.077)
GAIN			0.317***	0.068*
			(0.017)	(0.038)
DTV	0.042***	0.027***	0.031***	0.029***
	(0.005)	(0.005)	(0.005)	(0.005)
Predicted base price	1.017***	1.026***		
	(0.012)	(0.009)		
Index	0.443***	0.439***		
	(0.011)	(0.011)		
Predicted price			1.037***	1.032***
			(0.008)	(0.007)
Previous residual		0.519***		0.464***
		(0.017)		(0.047)
log(Holding time)	0.004*	−0.012***	0.062***	0.0001
	(0.002)	(0.002)	(0.006)	(0.008)
Constant	−0.535***	−0.551***	−0.823***	−0.496***
	(0.174)	(0.145)	(0.117)	(0.106)
N	31,966	31,055	32,044	30,450
Adj. R sq.	0.937	0.958	0.953	0.960
LOSS>0 (% of N)	6.657	6.743	6.856	7.087
VIF LOSS	1.054	1.175	1.121	1.178
VIF GAIN			1.871	3.792
F-stat.(LOSS=GAIN)			61.859	56.819
p-value(LOSS=GAIN)			<0.001	<0.001
LOSS-GAIN			0.607	0.421

Notes: The table presents results from replicating the first two columns of Table 2 in Genesove and Mayer (2001), and by adding prospective gains and using a hedonic price rather than a base price and the index, as in Bokhari and Geltner (2011). Note that the GM replication relies on different predictions than those in the two last columns; see Appendices A.1 and A.2 for more details. This is the reason for the different number of observations in the two model approaches. The dependent variable is log(List price). Predictions are of log(Selling price). Holding time is in weeks. Standard errors are calculated using a wild bootstrap (R=1,000) and clustered on list year and 3-digit zip codes. Significance: * p<0.1, ** p<0.05, *** p<0.01.

based on the backward-looking approach as explained in Appendix A.1. Additional information in the bottom of this table includes an F-test for the equality of the coefficient estimates on loss and gain, as well as the difference between these point estimates.

The results indicates that prospective gain may be an omitted variable in the first two columns, because the upper bound is more than halved when including the gain. The difference in coefficient estimates reflects the additional effect of facing a loss when determining the list price, interpreted in the same way as the estimates in the first two columns. Because the lagged residual is nested in both of the prospective variables, the variance inflation factor (hereafter VIF) is provided. The VIF of prospective gain is 3.79 in column (4), indicating that adding the residual control makes the variance of the gain more inflated. In addition to the direct effect from the residual, there may be an additional effect through the prospective loss. Yet, the relatively small VIFs indicate that using the residual control together with the gain does not impose a threat.[23] The purpose of including the residual is to offset the upward bias, intentionally including a control variable that shares common components with the variables of interest.

The results of estimating the relationship in equation (6) are presented in Table 4. The first column shows that the estimated effect is very small and no longer significant. To control for unobserved heterogeneity, the difference between the lagged log selling price and the lagged appraisal value is included as a control variable. This is intended to mimic the residual control variable used to downward bias the estimate in the hedonic-based models. Including this variable had little impact on the estimate. Columns (3) and (4) are results from estimating the hedonic-based model in equation (3) using the same observations as in the first two columns, meaning observations in the appraisal sample. The estimates are of a similar magnitude as those found when estimating the model specification on the full sample. For completeness, the hedonic-based model is estimated on a subsample of the hedonic sample that does not overlap with the appraisal sample period. Results from this exercise are provided in the two last columns, showing smaller estimates than for the pre-cutoff period. The lower estimates in the later period can be explained by several factors. For instance, there is likely a lower degree of autocorrelation in this subsample because the shorter time frame leads to fewer observations of the same units. Additionally, the predictions are more accurate during this later period. Another explanation is that the transition from using appraisal values to AVM predictions may have resulted in a structural change in the housing market: realtors provide the input in the form of information about the housing units when generating AVM predictions, potentially granting the realtors with more influence over seller expectations and the choice of list prices.

The results from estimating the appraisal-based model (6) suggest that loss aversion does not influence sellers in their choice of list prices. Yet, the coefficient estimate on prospective gain is significantly positive. A potential explanation of this is that these sellers are more likely to trade off some of the selling price for a faster sale. There is a clear difference between the estimates of the two models, both in magnitude and significance, which in light of the reduced-form investigation may be attributed to unobserved heterogeneity. This is investigated in more

[23]The typical rule of thumb thresholds for multicollinearity are a VIF above 5 or 10.

Table 4: Appraisal values as expected selling price

	(1)	(2)	(3)	(4)	(5)	(6)
LOSS	0.016	0.010	0.964***	0.609***	0.770***	0.398***
	(0.043)	(0.043)	(0.050)	(0.042)	(0.136)	(0.089)
GAIN	0.022***	0.021***	0.413***	0.181***	0.330***	0.066
	(0.003)	(0.003)	(0.019)	(0.032)	(0.030)	(0.052)
DTV	−0.003***	−0.003***	0.037***	0.035***	0.041***	0.037***
	(0.0005)	(0.001)	(0.006)	(0.006)	(0.005)	(0.007)
log(Appraisal value)	1.002***	1.003***				
	(0.002)	(0.002)				
Prev. log(SP)-log(AV)		0.002*				
		(0.001)				
Predicted Price			1.027***	1.023***	0.997***	1.012***
			(0.009)	(0.009)	(0.012)	(0.011)
Previous Residual				0.385***		0.492***
				(0.035)		(0.063)
log(Holding time)	0.005***	0.005***	0.055***	0.008	0.075***	0.001
	(0.001)	(0.001)	(0.005)	(0.005)	(0.008)	(0.015)
Constant	−0.075***	−0.082***	−0.636***	−0.385***	−0.262	−0.189
	(0.026)	(0.027)	(0.138)	(0.131)	(0.201)	(0.172)
Substitute for μ_{it}	P_{it}^{AV}	P_{it}^{AV}	$\hat{P}_{it}$	$\hat{P}_{it}$	$\hat{P}_{it}$	$\hat{P}_{it}$
Period	05Q1-16Q2	05Q1-16Q2	05Q1-16Q2	05Q1-16Q2	16Q3-20Q4	16Q3-20Q4
N	15,884	13,063	15,884	13,063	15,063	14,871
Adj. R sq.	0.996	0.996	0.950	0.956	0.929	0.943
LOSS>0 (% of N)	3.349	3.728	8.474	9.025	6.227	6.220
VIF LOSS	1.053	1.061	1.122	1.270	1.126	1.159
VIF GAIN	1.598	1.603	1.563	3.911	2.071	4.370
F-stat.(LOSS=GAIN)	0.020	0.078	112.345	138.147	7.661	30.009
p-value(LOSS=GAIN)	0.887	0.780	<0.001	<0.001	0.006	<0.001
LOSS-GAIN	-0.006	-0.012	0.550	0.429	0.440	0.332

Notes: Columns (1) and (2) use the appraisal values as substitutes for price expectations, using the appraisal sample, while (3) and (4) use hedonic predictions as substitutes estimated using the same sample. Columns (5) and (6) use post appraisal value cut-off observations from the hedonic sample, meaning from Q3 2016 to Q4 2020. The dependent variable is log(List price). Predictions are of log(Selling price). Holding time is in weeks. The market value in DTV is log of appraisal value in the appraisal-based models and the hedonic prediction of log selling price in the hedonic-based models. Standard errors are calculated using a wild bootstrap (R=1,000) and clustered on list year and 3-digit zip codes. Significance: * p<0.1, ** p<0.05, *** p<0.01.

detail in the next section.

4 Exploring the difference in estimates

4.1 Assessing unobserved heterogeneity

The difference in the seller loss aversion estimates is likely due to unobserved heterogeneity, because the predictions do not account for as much information as is inherently nested in the appraisal values. In what follows, I assess both time-invariant and time-varying unobserved heterogeneity that may cause the difference between the hedonic-based model estimates and the appraisal-based model estimates.

4.1.1 Constant and varying unobserved heterogeneity

There are mainly two types of unobserved heterogeneity, namely time-invariant and time-varying heterogeneity. Unobserved heterogeneity is assumed by GM to be only time-invariant. To investigate how this assumption affects interpretation, a closer investigation is conducted of the part of the residuals that originates from the mismeasurement of prospective variables. Denote the difference between the ideal and observed prospective loss as $\Delta\ell = LOSS^*_{ist} - LOSS_{ist}$. The three cases to be considered are:

$$\Delta\ell = \begin{cases} (\delta_s - \delta_t + w_{is})^+ - (\delta_s - \delta_t + v_i + w_{is})^+ & \text{if } v_{is} = v_{it} & \text{(i)} \\ (\delta_s - \delta_t + v_{is} - v_{it} + w_{is})^+ - (\delta_s - \delta_t + v_{is} + w_{is})^+ & \text{if } v_{is} \neq v_{it} & \text{(ii)} \\ (\delta_s - \delta_t + v_{is} - v_{it} + w_{is})^+ - (\delta_s - \delta_t + v_{is} - \tilde{v}_{it} + w_{is})^+ & \text{if } v_{is} \neq v_{it} \text{ and } P^{AV}_{it} \text{ for } \hat{\mu}_{it}. & \text{(iii)} \end{cases}$$

When moving from assuming constant unobserved heterogeneity (i) to allowing for time-varying unobserved heterogeneity (ii) a new source of losses emerge: changes in unobserved heterogeneity. Examples of time-varying unobserved heterogeneity include deterioration or renovation of the housing unit, noisy neighbors, and densification not captured at the aggregate level. Furthermore, if keeping the aggregate price level constant, renovating an apartment should increase the probability of facing a gain, while severe deteriorated should increase the probability of facing a loss. The change in quality is not captured in the observed prospective loss; it only captures the quality level at the previous sale.

In case (i), the unobserved heterogeneity makes the observed prospective loss include more noise than the (ideal) unobserved loss. If the unobserved heterogeneity is treated as random noise, this should result in attenuation. Case (ii) has fewer terms in the observed prospective loss compared to the unobserved prospective loss, resulting in fewer sources of variation in the observed loss. In case (iii), there should also be less variation in the observed loss, but this should be much closer to the true loss than in the other cases.

When unobserved heterogeneity is allowed to vary over time, it becomes evident that the observed prospective loss omits certain information due to the substitution with predictions.

The omitted information is not purely random within the model framework: it is an omitted variable. By following this notion, it is clear that even in the simple case (i), the measurement error is not causing attenuation. The potential consequences in case (i) are discussed by GM. However, biases arising from estimating the relationship should be considered in the context of both cases (ii) and (iii), particularly in terms of Berkson measurement error (Berkson, 1950). In a study examining the consequences of Berkson error, Haber et al. (2021) show how two types of Berkson error can introduce biases. The worst-case scenario is when the measurement error is differential, which can lead to spurious results.[24] However, when an explanatory variable suffers from Berkson error in form of random noise, meaning it is non-differential, this results in attenuation.[25] Differential measurement error within the model framework is discussed in more detail in Section 4.2.

Conversely, when using appraisal values, the unobserved heterogeneity should be smaller, because the surveyor observes both time-invariant and time-varying unobserved heterogeneity. If treating appraisal values as estimates, and the hedonic-based model suffers from differential measurement error, then the appraisal-based model should be preferred. By assuming constant unobserved quality, v_i, this implies disregarding a possible sources of why some sellers face losses. In turn, this can affect interpretation of how the measurement error in prospective variables affect estimates.

4.1.2 Unit fixed effects

Although the Oslo market is relatively small compared to other markets, there are still multiple repeated repeat purchases-to-listings in the sample, as presented in Table 2. This aspect of the sample is utilized to estimate the model while accounting for housing unit fixed effects, which should help limit concerns about time-invariant unobserved heterogeneity. Still, this could lead to issues with sample selection, because unobserved time-varying factors may cause sellers to hold their housing units for a relatively short duration. The two models in equations (3) and (6) are estimated for two panels each, namely panels of two-times and three-times repeated purchases-to-listings. Only units observed with two (three) repeated purchases-to-listings are included, making this exercise a test of robustness as well. Both pooling and unit fixed effects (hereafter FE) models are estimated for each balanced panel.

In Table 5, the pure hedonic-based model (3) and the pure appraisal-based model (6) are estimated using subsamples of two-times repeated observations, making the estimations equivalent to first differences. The hedonic unit FE estimates are smaller and more in line with GMs

[24]The two types are non-differential and differential Berkson errors. In the former, the residual from predicting P_{it} does not correlate with explanatory variables in the main specification, while in the latter they do.

[25]This can easily be seen from a simple measurement error example. Let $y = \beta x + \epsilon$ be the causal model, which we estimate as $y = \beta\tilde{x} + \tilde{\epsilon}$, with $\tilde{x} = x - u$ and $\tilde{\epsilon} = \epsilon + \beta u$. Then $\hat{\beta} \xrightarrow{p} \beta\sigma_x^2/(\sigma_x^2 + \sigma_u^2)$, in which σ denotes the standard deviation.

Table 5: Unit fixed effects - two-times repeated purchases-to-listings

	(1)	(2)	(3)	(4)
LOSS	0.423***	0.488***	−0.082	0.048
	(0.049)	(0.053)	(0.088)	(0.091)
GAIN	0.106***	0.151***	0.019***	0.032***
	(0.017)	(0.023)	(0.006)	(0.006)
Substitute for μ_{it}	$\hat{P}_{it}$	$\hat{P}_{it}$	P_{it}^{AV}	P_{it}^{AV}
Residual control	No	Yes	No	Yes
N	10,324	10,324	4,026	4,026
Adj. R sq.	0.937	0.939	0.976	0.977
LOSS>0 (% of N)	7.294	7.294	4.148	4.148
VIF LOSS	1.134	1.156	1.073	1.096
VIF GAIN	2.075	2.969	1.457	1.575
F-stat.(LOSS=GAIN)	35.186	39.697	1.351	0.034
p-value(LOSS=GAIN)	<0.001	<0.001	0.245	0.854
LOSS-GAIN	0.317	0.337	-0.101	0.016

Notes: The table presents results of unit fixed effect using balanced two-times repeated observations subsamples of purchases-to-listings. The dependent variable is log(List price). Predictions are of log(Selling price). The covariates include the log of holding time in weeks, DTV, the constant term for the pooling estimates, the predicted price, the log appraisal value, and the residual control (lagged selling price/appraisal value difference), which are omitted from the table. Standard errors are calculated using a wild bootstrap (R=1,000) and clustered on list year and 3-digit zip codes. Significance: * p<0.1, ** p<0.05, *** p<0.01.

results. The residual control does not bias the estimated effect downward. Yet, the smaller estimates indicate that there are unobserved constant heterogeneity affecting the estimates. Repeating the same exercise for the appraisal-based model still results in small and insignificant estimates of the effect. The pooling results from estimation on this subsample are in line with the baseline results in Table 3 (see Table D.4 in the Appendix). Unit FE estimations on a balanced panel with three-times repeated purchases-to-listings result in even smaller estimates (see Table D.5 in the Appendix). However, the sample size is small, especially for the appraisal-based variants, with relatively few cases of positive prospective losses.

Together with the appraisal-based estimates in Table 4, the results indicate that the estimated effect is biased by unobserved heterogeneity. The unit fixed effects does not capture time-varying unobserved heterogeneity, which could be the reason why results in columns (1) and (2) still show significant results. Adding to this notion, when measuring the prospective variables as the difference in aggregate price levels, with prospective loss defined as $LOSS_{ist} = (\delta_s - \delta_t)^+$, the results indicate that changes in aggregate price levels are important for list prices when estimating the hedonic-based model. However, the effect disappears when estimating the appraisal-based model. The price index, which is constructed from a linear hedonic model, could make the list price model suffer from Berkson residuals. Ultimately, adding the lagged residual control may bias the estimated effect upwards. Results are presented in Table D.6 in the Appendix.

4.1.3 Some time-varying unobserved heterogeneity in predictions

To establish whether there is a problem caused by unobserved heterogeneity in the hedonic model estimations, the residual obtained from fitting a hedonic model of the appraisal values is included as an explanatory variable in the hedonic model of the log selling price. This mimics same procedure presented by Anundsen and Røed Larsen (2018, p. 2144). This implies including $e_{it}^{AV} = P_{it}^{AV} - \hat{P}_{it}^{AV}$, and then estimating the list price model with predictions $\hat{P}_{it}|_{e_{it}^{AV}} = X_i\hat{\beta} + \hat{\delta}_t + \hat{\beta}_e e_{it}^{AV}$ as substitutes. Effectively, e_{it}^{AV} should be a compound factors that the econometrician does not observe but the surveyor does. As such, this residual does not include the difference between the actual selling price and the price expectation (w_{it}). This gives the following feasible model:

$$L_{ist} = \alpha_0 + \alpha_1 \hat{P}_{it}|_{e_{it}^{AV}} + m_l(P_{is} - \hat{P}_{it}|_{e_{it}^{AV}})^+ + m_g(P_{is} - \hat{P}_{it}|_{e_{it}^{AV}})^- + \eta_{it}. \tag{7}$$

By following this three-step procedure, this ensures that if the residuals from the appraisal value estimation are irrelevant for the selling prices, they would not add more information to the predictions thus not affecting the loss aversion estimates.

The results from estimating relationship (7) are presented in Table 6. In columns (1) and (2), the loss aversion effect is somewhat smaller compared to the baseline hedonic-based estimates using the same sample period (see Table 4). Using unit fixed effects with these alternative predictions in columns (3) and (4) yields even smaller estimates.[26] This indicates that the appraisal value residual adds more information to the predictions, so that prospective loss and gain take account for more of the time-varying unobserved heterogeneity.

4.1.4 Choice of dependent variable in the hedonic model

As seen above, estimating the appraisal-based model yielded insignificant results. This gives rise to the question about whether the hedonic-based results would be different if the dependent variable in the hedonic price model was the appraisal value. Meaning, whether predictions should be of the appraisal values. Modeling price expectations on selling prices implicitly assumes that the expectations follow the same movements and distributional properties of those prices. In markets where experts and their opinions are valued, the price expectations could

[26] Pooling results on the same observations as in columns (3) and (4) yielded estimates of 0.54 and 0.47, respectively, both significant at the 1 percent level.

Table 6: Unobserved heterogeneity through appraisal value residuals

	Eq. (7)		Eq. (7), unit FE		Eq. (8)	
	(1)	(2)	(3)	(4)	(5)	(6)
LOSS	0.491***	0.480***	0.314***	0.267***	0.889***	0.039***
	(0.083)	(0.082)	(0.086)	(0.085)	(0.035)	(0.009)
GAIN	0.071***	0.100***	0.054***	0.076***	0.396***	0.019***
	(0.012)	(0.017)	(0.007)	(0.011)	(0.019)	(0.003)
Residual control	No	Yes	No	Yes	No	Yes
N	15,153	11,936	4,264	3,132	15,949	15,949
Adj. R sq.	0.990	0.991	0.966	0.962	0.950	0.996
LOSS>0 (% of N)	2.831	3.050	3.612	4.119	12.283	12.283
VIF LOSS	1.044	1.055	1.090	1.116	1.146	1.272
VIF GAIN	1.617	1.694	1.521	1.711	1.573	2.131
F-stat.(LOSS=GAIN)	22.574	18.122	8.784	4.865	157.259	4.739
p-value(LOSS=GAIN)	<0.001	<0.001	0.003	0.027	<0.001	0.030
LOSS-GAIN	0.419	0.380	0.259	0.191	0.493	0.020

Notes: The table presents results using hedonic predictions with appraisal value residuals as a control variable, and using hedonic predictions of appraisal values as substitutes. The unit fixed effects use a balanced two-times repeat purchases-to-listings subsample. The dependent variable is log(List price). The covariates include the log of holding time in weeks, DTV, the price prediction, and the residual control, which are omitted from the table. Note that the unit fixed effects are conducted using demeaning, making the share of positive losses potentially misleading. Standard errors are calculated using a wild bootstrap (R=1,000) and clustered on list year and 3-digit zip codes. Significance: * p<0.1, ** p<0.05, *** p<0.01.

be affected by these third-party agents. In the Norwegian setting, people listen to realtors and surveyors, potentially anchoring to their suggestions of list price setting. Moreover, if potential buyers observe the appraisal value as well, this may constrain sellers who want to set list prices above appraisal values.[27] Thus, the appraisal value may be the better option.

By using the same hedonic specification as for the model that produced the predictions of selling price, the importance of using fitted values of appraisal values as substitutes for price expectations is assessed by estimating the following model:

$$L_{ist} = \alpha_0 + \alpha_1 \hat{P}_{it}^{AV} + m_l(P_{is} - \hat{P}_{it}^{AV})^+ + m_g(P_{is} - \hat{P}_{it}^{AV})^- + \eta_{it}, \tag{8}$$

in which $\hat{P}_{it}^{AV}$ is the fitted value of log appraisal value. The results are presented in columns (5) and (6) in Table 6. Column (5) shows that the choice of dependent variable in this case does not matter for the estimates. To bias the estimates downward in column (6), the non-lagged

[27] There might also be anchoring effects of sales that are close in distance in multiple dimensions, such as in time, space, and in characteristics of the housing unit.

residual is used as a control variable. This is done because the list price is decided upon after receiving the appraisal value. This residual is the same as used in the hedonic predictions in the first two columns, so that it captures what the surveyor observes but not the econometrician, making it a noisy proxy for both time-invariant and time-varying unobserved heterogeneity. Adding this makes the effect almost disappear, yet, the estimate is significant at the 5 percent level. The two last columns are crucial for understanding the biases that arise when estimating the relationship. While there is a strong upward bias when not including the residual control, adding the residual leaves a very small effect.[28]

4.2 Differential measurement error

The results show that by dealing with some of the unobserved heterogeneity the estimates become smaller. Specifically, the unit fixed effects estimates only leaves the time-varying unobserved heterogeneity in the housing units out, which yielded smaller estimates. When providing even more information to the hedonic models this reduced the estimates as well. The combination of unit fixed effects and better hedonic predictions yielded to an even smaller estimated effect.

In light of these results, I propose a different perspective for considering the biases that arise when estimating the relationship in (3). That being, as briefly mentioned in Subsection 4.1.1, that the error structure is called *differential (Berkson) measurement error*. The purpose of this perspective is to provide insight into the complexity of the magnitude and direction of the bias. The measurement error is differential when the residual from the prediction of selling price is correlated with the dependent variable (Haber et al., 2021, p. 865). This means that the residuals from the hedonic predictions are correlated with the prospective loss and gain, or other control variables, while being relevant for the list price. It is clear from the decomposition in Appendix B.1 that the residual $v_{it} + w_{it}$ from predicting $\hat{P}_{it}$ should be correlated with both $LOSS_{ist}$ and $GAIN_{ist}$ due to the part $v_{is} + w_{is}$. By letting the part of the residual w_{is}, w_{it} be identical and independently distributed, the correlation between residuals is due to the correlation between the unobserved qualities v_{is} and v_{it}. Further, $LOSS_{ist}$ and $GAIN_{ist}$ suffer from differential measurement error because they are constructed by the predictions as well. The prediction residual from the hedonic model is referred to as the Berkson residual.

The reduced-form framework presented in Appendix B.1 is a simple way of determining potential biases, but the biases become less tractable when allowing unobserved heterogeneity to vary over time. Moreover, omitted variables in the list price model that correlate with the Berkson residual could complicate inference even further. This is why it is fruitful to investigate the biases more rigorously, drawing from the discussion by Haber et al. (2021). A simplified case is by estimating the list price model following the approach by GM of using $LOSS_{ist}$ alone

[28] In equation (B.7) in the Appendix, adding the non-lagged residual approximately eliminates $\alpha_1 v_{it}$.

(omitting prospective gains). For simplicity, leave out the subscripts. Assume that $\hat{P}$ is observed instead of μ, in which $\mu = \hat{P} + e$, and let $LOSS$ be the mismeasured $LOSS^*$. Let $LOSS$ be the prospective loss in the causal relationship, which is a simplification explained in further detail in Appendix B.2. The causal relationship is given by $L = \lambda_0 + \lambda_1 \mu + \lambda_l LOSS + \lambda_\pi \pi + \tilde{\eta}$, with $\tilde{\eta}$ being noise with standard normal distribution.

Denote the covariance between μ and prospective loss as $Cov(\mu, LOSS) = \sigma_{\mu,l}$, the variance as $Var(LOSS) = \sigma_l^2$, and assume $\sigma_{e,l} \neq 0$ (> 0), $\sigma_{\mu,l} \neq 0$ (< 0). Let neither of $\{LOSS, \mu\}$ correlate with the omitted variable(s) in the list price model. Assuming no correlation means that the true α_1 is estimated even when the causal relationship between list price and the other variables is not perfectly modeled. Meaning, assume that there is no omitted variable bias in estimating λ_1, so that $E(\alpha_1) = \lambda_1$.

Further, omit one variable π when estimating the model. Let the list price residual $\eta = \lambda_\pi \pi + \tilde{\eta}$ with $\sigma_{\pi,e} \neq 0$ and $\sigma_{\tilde{\eta},e} = 0$, but for simplicity keep assuming $\sigma_{\pi,\mu} = \sigma_{\pi,l} = 0$. Ultimately, this leaves two sources of which the Berkson error, e, can affect the estimate of m_l: through $\hat{P}$ and through the omitted variable π. Then, it can be showed, see Appendix B.2, that the bias of estimating the relationship $L = \alpha_0 + \alpha_1 \hat{P} + m_l LOSS + \eta$ has the following form:[29]

$$E\left(\hat{m}_l(\hat{P}) - \hat{m}_l(\mu)\right) = \lambda_\pi \frac{\sigma_{\pi,e}(\sigma_{\mu,l} - \sigma_{e,l})}{\zeta} + \alpha_1 \frac{\sigma_{e,l}}{\sigma_l^2} \frac{(\sigma_\mu^2 - \sigma_e^2)\sigma_l^2}{\zeta}, \tag{9}$$

in which

$$\zeta = (\sigma_\mu^2 - \sigma_e^2)\sigma_l^2 - (\sigma_{\mu,l} - \sigma_{e,l})^2.$$

λ_π is the true coefficient on the omitted variable π. The bias consists of two parts, the first is the bias from the correlation between the unobserved variable π in the list price model and the Berkson residual. The second is the bias from the correlation between the prospective loss variable and the Berkson residual. If these correlations are set to zero there is no bias. Also, the second part will disappear if the price prediction does not affect the list price ($\alpha_1 = 0$), or if the expected market price was perfectly observed ($\sigma_e^2 = 0$ giving $\sigma_{e,l} = 0$). Without an omitted variable correlating with the Berkson residual, there would still be bias. By construction, $\sigma_\mu^2 = \sigma_{\hat{P}}^2 + \sigma_e^2$, implying $\sigma_\mu^2 > \sigma_e^2$ if some variation in μ is accounted for in $\hat{P}$. Thus, the denominator is non-negative if there is a correlation between the prediction and the prospective loss, $\zeta \geq 0$.[30]

The sign on the second part depends on the covariance between the Berkson error, e, and $LOSS$, and the denominator, ζ, of the last fraction, which is non-negative. Because the expected market price enters $LOSS$ with a negative sign, it is likely that $\sigma_{\mu,l} < 0$ while recalling that $\sigma_{e,l} > 0$, which is due to the common time-invariant component, ω_i. Further, note that

[29] For assumptions regarding the predictions of $\hat{P}$, confer with Haber et al. (2021) Section 2.

[30] Note that $\sigma_{\hat{P}}^2 = \sigma_\mu^2 - \sigma_e^2$ and $\sigma_{\hat{P},l} = \sigma_{\mu,l} - \sigma_{e,l}$, so that $\zeta = \sigma_{\hat{P}}^2 \sigma_l^2 - \sigma_{\hat{P},l}^2 \geq 0$.

the difference between these are the covariance between the prediction and the loss, $\sigma_{\hat{P},l} \leq 0$. Therefore, the second part is likely to contribute with a downward bias in the estimate of m_l.

Considering the first part of equation (9), in general, one can expect omitted variables in the list price model to correlate with the residual in the price prediction because there could be unobserved common factors that determine both. Even if these are not correlated with any other covariates in the list price model, the bias still enters with an unknown sign. If $\sigma_{\hat{P},l} < 0$ and $sign(\lambda_\pi) = sign(\sigma_{\pi,e})$, this bias is negative. There could still be unobserved heterogeneity, so that $sign(\lambda_\pi) \neq sign(\sigma_{\pi,e})$, yielding a positive bias. This bias remains even when there is no correlation with the prospective loss and the Berkson error.

In order to say something more about the unobserved heterogeneity, Table D.3 in the Appendix shows the shares of observations having a positive or negative sign at the two points in time $\{s, t\}$, split into those not facing and those facing an appraisal-based prospective loss. For those with positive loss, there is a higher share of under-valuation from the hedonic model in both periods. There is also a higher share of under-valuation in the first period and over-valuation in the second period for the positive losses. Whether under-valuation comes from positive unobserved heterogeneity or a random over-payment is hard to pinpoint. When regressing an indicator of positive loss on indicators of whether the sign on the residuals are positive, this yields the average marginal effects presented in Table 7. On average, an under-valuation in s increases the probability of the appraisal-based loss being positive by 6.5 percent, while under-valuation in t decreases the probability by 1.6 percent. On the other hand, the hedonic-based prospective loss is more affected the first residual but a opposite sign from the second. If disregarding w_{is}, the latter could indicate a larger common time-invariant part, ω_i, positively determining the loss. The different signs of the former could indicate that some of the unobserved heterogeneity is accounted for.

Table 7: AME for indicator of positive appraisal-based loss on hedonic model residual signs

	(A) $\mathbb{1}\{LOSS^{AV}_{ist} > 0\}$			(B) $\mathbb{1}\{LOSS_{ist} > 0\}$		
	AME	SE	p-value	AME	SE	p-value
$\mathbb{1}\{e_{is} > 0\}$	0.065	0.007	<0.001	0.184	0.008	<0.001
$\mathbb{1}\{e_{it} > 0\}$	-0.016	0.004	<0.001	0.080	0.005	<0.001

Notes: The table reports the average marginal effects (AME) from logistic regression. The dependent variable is an indicator for whether the seller faces a positive loss, and this is regressed on indicators for the signs of the residuals. All regressions include an interaction term between the two indicators. In the prospective loss, μ_{it} is substituted with P^{AV}_{it} in panel A and with $\hat{P}_{it}$ in panel B. Standard errors are heteroskedasticity-robust. N=14,544 in panel A and N=30,138 in panel B.

In this simplified exercise, when assuming that the prospective terms are correctly specified and setting many correlations to zero, having a control variable with differential measurement error and an omitted variable correlating with this error may bias the estimate. The two terms in equation (9) are seemingly independent of λ_l, so even if the true effect is zero ($\lambda_l = 0$), estimates can still be significant.[31] There would still be bias even if there was no correlation between loss and the true market price μ. Ultimately, this reduces to a *spillover effect*. Further, by letting the measurement error in $\hat{P}$ be present in the prospective loss, the total effect would be much harder to trace and interpret.

Adding the residual from the hedonic price prediction, $v_{is} + w_{is}$, further complicates interpretation. The bias expression now becomes even larger as the residual likely correlates with the other variables, including the omitted variable, π, so that there is no obvious way to interpret the sign on the bias expression (see Appendix B.2 for the full expression). Previous studies, such as GM, show that including the residual gives lower estimates than an estimation without the residuals, but whether a model with residuals always suffer from upward or downward bias is not clear. This is particularly the case when accounting for correlated unobservables in both the list price model and the hedonic model.

When applying this Berkson perspective to the first two columns in Table 6, the bias is smaller when the residual contains less factors correlating with the list price. Utilizing more information in the predictions therefore mitigate the bias. For the two last columns, the strong effect found by using fitted appraisal values iterates a point made by Haber et al. (2021): using predictions instead of observed variables may inflate coefficient estimates.

4.2.1 Establishing the source(s) of the biases

Recall specification (5). For convenience, denote the substitute of μ_{it} that is not part of the prospective terms as the *outside price*, while the substitute within the prospective terms is referred to as the *inside price*:

$$L_{ist} = \alpha_0 + \alpha_1 \underbrace{\mu_{it}}_{\text{outside price}} + m_l(P_{is} - \underbrace{\mu_{it}}_{\text{inside price}})^+ + m_g(P_{is} - \underbrace{\mu_{it}}_{\text{inside price}})^- + \epsilon_{it}. \tag{10}$$

To establish whether the biases originate from measurement error in the prospective terms, a *spillover effect* from the outside price, or a combination of both, the analysis alternates between the following substitutes for the outside and inside prices: the hedonic log selling price predictions ($\hat{P}_{it}$), the log appraisal values (P_{it}^{AV}), and the hedonic log appraisal value estimates ($\hat{P}_{it}^{AV}$). This breaks up the common source of measurement error between the main list price determinant, i.e, the outside price, and the prospective terms.[32]

Table 8 presents results from estimating the main model with different substitutes for the

[31] The coefficient that partly determines the size of the bias is the coefficient on expected market price, which should be close to 1.

[32] The procedure relates to the interpretation of potential biases as presented in equation (9), assessing the spillover effect of having measurement error in the outside price and an omitted variable.

Table 8: Mixing price substitutes

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
LOSS	-0.005	-0.007	0.728^{***}	-0.085	0.008	0.008	0.252	0.028
	(0.009)	(0.010)	(0.242)	(0.344)	(0.008)	(0.006)	(0.170)	(0.042)
GAIN	-0.004	-0.005	-0.134^{***}	-0.487^{***}	0.004^{*}	0.003	-0.131^{***}	0.020^{***}
	(0.003)	(0.003)	(0.016)	(0.031)	(0.002)	(0.002)	(0.010)	(0.003)
Outside price	P_{it}^{AV}	P_{it}^{AV}	$\hat{P}_{it}$	$\hat{P}_{it}$	P_{it}^{AV}	P_{it}^{AV}	$\hat{P}_{it}^{AV}$	$\hat{P}_{it}^{AV}$
Inside price	$\hat{P}_{it}$	$\hat{P}_{it}$	P_{it}^{AV}	P_{it}^{AV}	$\hat{P}_{it}^{AV}$	$\hat{P}_{it}^{AV}$	P_{it}^{AV}	P_{it}^{AV}
Residual control	No	Yes	No	Yes	No	Yes	No	Yes
N	15,884	14,185	15,884	14,573	15,949	14,240	15,949	15,949
Adj. R sq.	0.996	0.996	0.916	0.970	0.996	0.996	0.921	0.996
LOSS>0 (% of N)	8.474	8.530	3.349	3.603	12.283	12.556	3.373	3.373
VIF LOSS	1.125	1.126	1.054	1.062	1.157	1.161	1.054	1.054
VIF GAIN	1.527	1.571	1.576	1.832	1.540	1.590	1.573	1.616
F-stat.(LOSS=GAIN)	0.034	0.054	13.360	1.511	0.197	0.432	5.072	0.036
p-value(LOSS=GAIN)	0.854	0.817	<0.001	0.219	0.657	0.511	0.024	0.849
LOSS-GAIN	-0.001	-0.002	0.863	0.402	0.004	0.005	0.383	0.008

Notes: The table presents results of mixing the substitutes of outside and inside prices. The dependent variable is log(List price). The first four columns use hedonic predictions of log selling prices, while the last four columns use hedonic fitted values of appraisal values. This is reported in the rows below the point estimates. The covariates include the log of holding time in weeks, DTV, the constant term, the outside price, and the residual control, which are omitted from the table. Standard errors are calculated using a wild bootstrap (R=1,000) and clustered on list year and 3-digit zip codes. Significance: * p<0.1, ** p<0.05, *** p<0.01.

outside and inside prices. The choice of control variables depends on the outside price substitute. For instance, if the outside price is substituted with the hedonic prediction of log selling price, the GM residual control variable $(v_{if} + w_{if})$ is used, as in Table 3. The other downward-biasing controls are the same as those used in the estimations that gave the results in Table 6.

Most importantly, the results show that even when modeling the relationship with the appraisal-based loss and gain, using hedonic predictions as substitute for the outside price yields a large significant estimate of the upper bounds. Specifically, substituting with the hedonic predictions of selling price yield the highest estimates, while using the fitted appraisal values yield lower estimates with a insignificant lower bound. When using appraisal values as the outside price and hedonic selling price predictions as the inside price, the estimate becomes very small and insignificant. The difference in how the hedonic selling price predictions and the hedonic fitted appraisal values affect the estimates suggests that additional factors within the selling price residuals contribute to a stronger bias. In light of the institutional context, and the results from using fitted appraisal values in Table 6 which also suffer from differential Berkson error, substituting with appraisal values may be preferred.

These results suggest that the spillover effect from the main list price determinant is what mostly drives the difference between the hedonic-based and appraisal-based estimates. Thus, the choice of substitute for the main list price determinant (outside price) is a key factor for identification.[33]

4.2.2 Consequences of the spillover effect

As is now established, there is information in appraisal values besides what is already captured by hedonic prediction that may be beneficial for inference. Specifically, there seems to be a spillover effect from measurement error in the main list price determinant, that being the outside price above. To better understand how the spillover effect influences the loss aversion estimate and the importance of additional information in appraisal values, appraisal value residuals are added to the hedonic predictions by adjusting the weight assigned on these residuals. The procedure is intended to demonstrate how additional information in the main list price determinant results in lower and less biased estimates. Specifically, the procedure assigns weights, $\omega \in [0,1]$, on the appraisal value residuals and adds them to the predicted price, resulting in a new compound price $\tilde{P}_{it} = \hat{P}_{it} + \omega\tilde{v}_{it}$, in which $\tilde{v}_{it} = P_{it}^{AV} - \hat{P}_{it}^{AV}$. The model incorporates this to the main list price determinant but not in the construction of prospective losses and gains:

$$L_{ist} = \alpha_0 + \alpha_1 \tilde{P}_{it} + m_l(P_{is} - \mu_{it})^+ + m_g(P_{is} - \mu_{it})^- + \eta_{it}. \tag{11}$$

This procedure conducted for both of the two substitutes for μ_{it} when constructing the prospective terms, namely the hedonic predictions and the appraisal values. Results from estimating the relationship with the usual controls (DTV, log of holding time, a constant term) but without the residual control are presented in Figure 2.

This procedure accounts for different degrees of missing information and its impact on the loss aversion estimate. The coefficient on the main list price determinant, which is here substituted with $\tilde{P}_{it}$, is close to one. Therefore, adding the residual into this price ensures that the residual is weighted by approximately ω in the list price model. When adding the residual as a separate explanatory variable, the weight would be the coefficient, which is represented in the figure with the green dashed lines. As is clear from both panels, the estimated effect becomes smaller when more of the unobserved heterogeneity is included in the compound price. The influence is stronger in the hedonic-based loss and gain model (Figure 2 panel A) than in the appraisal-based loss and gain model (Figure 2 panel B). Assigning $\omega = 1$ still yields a significant estimate.

The next procedure introduces random noise rather than relevant information to further illustrate how the spillover effect influences estimates. The main list price determinant takes the form $\tilde{P}_{it} = \hat{P}_{it} + \xi$, with the normally distributed random noise $\xi \sim N(0, \sigma^2)$. The relationship in equation (11) is estimated over a range of $\sigma \in [0,1]$. Because the predicted price is in log

[33] This is at least the case for the sample containing relatively few positive prospective losses. The situation might differ with a larger share of listings facing prospective losses.

Figure 2: Effects of weighting in the appraisal value residual

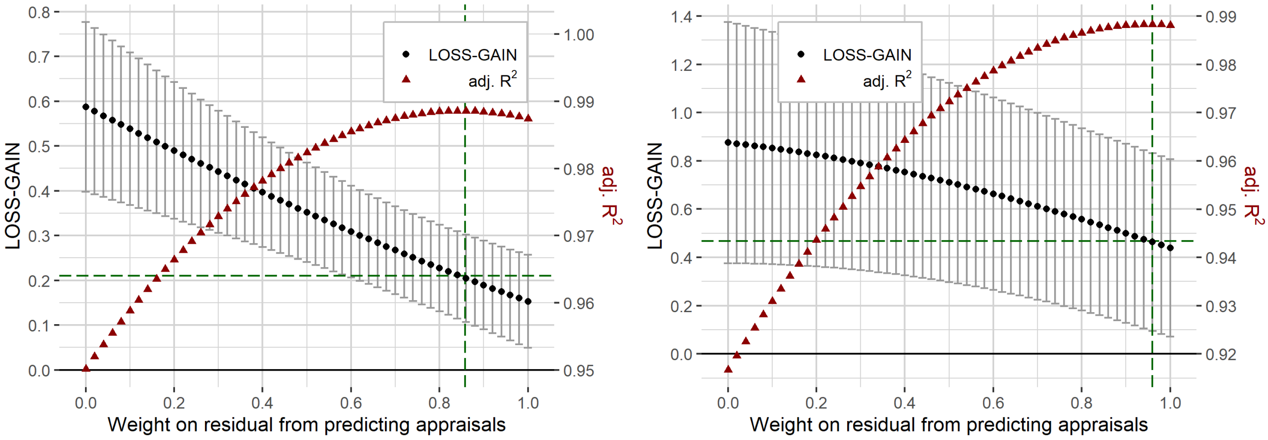

(A) Hedonic-based LOSS/GAIN (B) Appraisal-based LOSS/GAIN

Notes: The figure presents the difference between prospective loss and gain, a 95 percent confidence interval, and the adjusted R^2 for different weights, ω, in the compound price $\tilde{P}_{it} = \hat{P}_{it} + \omega \tilde{v}_{it}$, in which $\tilde{v}_{it} = P^{AV}_{it} - \hat{P}^{AV}_{it}$. Panel A shows the hedonic-based loss and gain difference, meaning using the hedonic price prediction as the inside price substitute. Panel B show the appraisal-based loss and gain difference. The green dashed lines represent the coefficient estimates on the residual (x-axis) and the effect (left y-axis) when estimating a model keeping the residual as a separate explanatory variable. Standard errors are clustered on list year and 3-digit zip codes.

scale, σ is capped at 1, so that even a $\xi \sim N(0,1)$ introduces a relatively large amount of random noise to the predictions.

Figure 3 tells two stories.[34] When substituting the main list price determinant with hedonic predictions, adding more noise to the hedonic outside substitute increases the estimated effect while also increasing the confidence interval. When substituting the main list price determinant with appraisal values, the effect becomes smaller: starting off with a positive and significant estimate and ending up with a negative significant significant. By inflating the Berkson residual with random noise, the direction of the bias depends on the construction of the prospective losses and gains. The spillover effect is particularly detrimental to the hedonic-based model, suggesting that poor predictions may result in higher estimates, even when the added noise is random.[35]

[34]Figure D.4 in the Appendix adds the lagged residual as a control variable to account for unobserved heterogeneity and potentially downward bias the estimates, showing the same pattern as in Figure 3.

[35]The reduced-form presentation of the model as presented in Appendix B.1 is not able to explain these results, i.e., in equation (B.8), the new residual becomes $u_{it} - \alpha_1 \xi$, in which ξ is independent and identically distributed.

Figure 3: Effects of adding random noise, without lagged residual

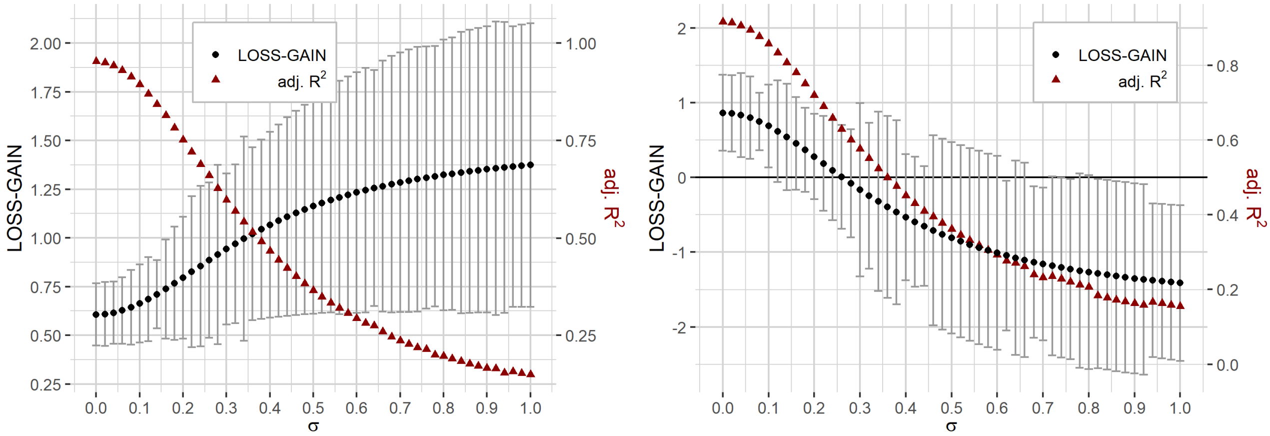

(A) Hedonic-based LOSS/GAIN (B) Appraisal-based LOSS/GAIN

Notes: The figure presents the difference between prospective loss and gain, a 95 percent confidence interval, and the adjusted R^2 when adding noise $\xi \sim N(0, \sigma^2)$ with different standard deviations σ to the predicted price $\hat{P}_{it}$. For each $\sigma > 0$, the procedure of drawing ξ and estimating the coefficient difference is repeated 1,001 times and the median coefficient difference between loss and gain, with the corresponding confidence interval, is chosen to be plotted. This is done because the loss-gain coefficient distribution becomes normal when adding normally distributed noise. Panel A shows the hedonic-based loss and gain difference, meaning using the hedonic price prediction as the main list price determinant substitute. Panel B show the appraisal-based loss and gain difference. Standard errors are clustered on list year and 3-digit zip codes.

4.3 Discussion

The results show that both time-invariant and time-varying unobserved heterogeneity in prices are important sources of biasing the estimates when substituting true expectations with predictions of selling prices. Dealing with the time-invariant unobserved heterogeneity supports the argument that this leads to upward bias in estimates, as discussed by Genesove and Mayer (2001). Yet, when going beyond their treatment of the subject matter, by investigating the role of time-varying unobserved heterogeneity through different exercises, it is clear that this is also an important biasing factor, even in the proposed lower bound estimate of the effect. Feasible inference requires dealing with the differential measurement error in the price expectation substitute, which is possible by retrieving better substitutes for these expectations. An example of a better substitute is the appraisal value.

A question that still remains to be answered is whether sellers in the Norwegian housing market are loss averse. The result from estimating relationship (6) suggests that sellers are not nominally loss averse, but inference relies on how likely it is that the estimates do not suffer from omitted variable bias. As discussed in Subsection 3.3, there may still be unobserved heterogeneity, represented by the surveyor error, and the direction of the bias in this context is probably intractable.

The institutional context in which the sellers are operating in also matters in understanding the estimate. Sellers are presented with appraisal values by the surveyors, and these surveyors are paid by the sellers to do so. Meaning, the surveyors are paid experts and their opinions should matter for the sellers. Moreover, the appraisal values serve the purpose of being suggestions for list prices. Therefore, appraisal values will likely serve as anchors for the sellers' expectations about prices. In addition, the sellers know that these appraisal values will be observed by potential buyers. Although not a binding constraint, the sellers may still feel constrained in their list price choices, because setting the list prices above the appraisal values will signal buyers about the type of these sellers. This in turn may not benefit them. Such a soft constraint may overpower the urge to act according to their impulses caused by being loss averse.

Overall, there are three take-home messages. First, using hedonic predictions in the list price model could introduce intractable biases. Second, prediction-based list price models are more sensitive to poorly specified hedonic models, which are used to produce predictions that are substituted into the list price models. Third, the appraisal-based model estimate is close to zero and not significant, a result that must be considered in light of the institutional setting. In the counterfactual case where buyers do not observed appraisal values, this may allow the sellers to exhibit loss averse behavior. This institutional dependence emphasizes the difference between being loss averse and acting in accordance with loss aversion.

5 Robustness and sensitivity

5.1 More on the importance of time-varying unobserved heterogeneity

The results highlight the potential for time-varying unobserved heterogeneity to contribute in biasing the estimates. Another test of how sensitive the results are to time-varying heterogeneity is to separate the residual from hedonic estimation of appraisal values into a time-invariant and time-varying part. The residual represents the unobserved heterogeneity in appraisal values, and this is separated into two parts: an explained and an unexplainable part. The explained part consists of the fitted values from regressing the appraisal value residuals at time t on the appraisal value residuals at time s, while also controlling for different time-varying measures, including local appreciation measured by a repeated sales index.[36] The unexplained part consists of the residuals from this regression, which should mostly nest time-varying heterogeneity. Controlling for the unexplained (time-varying) part in the hedonic-based list price model (3)

[36] Estimating this model resulted in an adjusted R^2 of 0.375.

resulted in a much lower estimate (see Table C.4). In contrast, controlling for the explained (mostly time-invariant) part of the residual did not reduce the estimate by much. This provides further evidence of the importance of the time-varying unobserved heterogeneity as a source of bias in the estimates (see Appendix C.2 for a more details).

5.2 Endowment effect or unobserved heterogeneity

Another approach to assess time-varying unobserved heterogeneity is by following Graddy et al. (2022) in their strategy of identifying the endowment effect. They study loss aversion and anchoring in the art market, and consider two subsamples defined by different lengths of holding time to identify the endowment effect. Sellers are affected by the endowment effect if they are exhibiting stronger loss averse behavior when the holding time is longer (Strahilevitz & Loewenstein, 1998). Graddy et al. (2022) find evidence of the endowment effect.

When performing a similar exercise for the housing market, longer holding times are likely to increase the presence of time-varying unobserved heterogeneity in the prospective terms. If the estimates from hedonic-based models are biased due to unobserved heterogeneity, the estimates should be larger when holding times are longer. The results are consistent with this, suggesting that there are time-varying unobserved heterogeneity that the surveyors might account for in their appraisal values (see Table D.7).[37]

5.3 Surveyors being loss averse on sellers' behalf

Surveyors are certified professionals, who may be offended if being accused of adjusting their appraisal values according to sellers' preferences, which goes against the surveyors' own code of conduct (Norsk takst, 2022). Yet, there might be some surveyors that are biased towards the sellers' preferences. Therefore, appraisal values may be influenced by sellers' loss aversion. When substituting hedonic predictions for price expectations (the hedonic-based model), the results suggest that surveyors suffer from loss aversion at the same extent as sellers. The estimated effect is close to the results presented in Table 4, columns (3) and (4). It is improbable that sellers influence surveyors to such a degree, which indicates that the Berkson error is biasing the estimates (see Table D.8).

[37] When investigating the difference between the lagged and non-lagged hedonic residuals, the sample variance of this difference is larger in the subsample comprising observations with longer holding times. This is an indicator for more unobserved heterogeneity being present in this subsample. Moreover, the probability of facing a loss decreases when holding time increases due to appreciation in aggregate prices. Estimating the appraisal-based model yielded an insignificant estimate when using the shorter holding time subsample, while the estimate was negative and significant at the 5 percent level when using the longer holding time subsample.

5.4 No trimming

To address whether there are issues with the data trimming that creates the difference between the hedonic-based and the appraisal-based estimates, the main models are estimated using a non-trimmed data. The observations that are dropped are the top influencing observations based on the Cook's distance in the baseline hedonic-based and appraisal-based model estimations. The results are on the same level of magnitude as those presented above, suggesting that the trimming is not causing the difference in estimates (see Table D.9 and Figure D.5).

5.5 Hockey stick

Prospect theory implies a kink at the transition point from prospective losses to gains. This should appear when comparing prospective losses and gains with the list price premium, which is defined as the list price to price expectation spread. When adding prospective losses and gains together, resulting in a variable that can be characterized as a reference dependence variable, a hockey stick pattern should emerge with a kink at zero when plotted against the list price premium. The hedonic-based prospective terms display a pattern resembling a hockey stick, while the appraisal-based prospective terms do not (see Figure D.6). Furthermore, the fitted appraisal-based relationship also displays a hockey stick-like pattern. However, there are no clear kinks at zero in the patterns displayed by the prediction-based variables, and the curve looks linear for prospective losses and convex for prospective gains. Adding to this, the models are estimated with the quadratic and cubic functional forms of the prospective gain, yielding insignificant estimates close to zero. Together, this suggests the differential measurement error in the prediction-based models results in estimates consistent with reference dependence but not loss aversion. Even when disregarding the observed appraisal values, using a model with linear curves with different slopes above and below the cutoff is ultimately coercing simple linear fits, while a more suitable model of the relationship includes a quadratic or cubic functional form of a non-truncated reference dependence variable (see Table D.10, in which the last columns include the reference dependence variable instead of the prospective gain).

6 Conclusion

The established literature about sellers' loss aversion in housing markets documents a strong and significant effect. A key variable for estimating the effect, which is used to construct the variables of interest and as a control variable, is the sellers' price expectations. This paper estimates the effect using repeat purchases-to-listings data from the Oslo housing market, employing a proposed more appropriate substitute for expectations than has been previously used: appraisal values. These appraisal values are provided by surveyors, which should better proxy true price expectations than standard hedonic predictions because they account for more of the heterogeneity unobserved by the econometrician, may influence sellers' expectations, and account for institutional factors.

To emphasize the potential improvement in estimates, the effect is estimated by first replicating previous studies (Bokhari & Geltner, 2011; Genesove & Mayer, 2001), resulting in similar strong and significant estimates. However, when estimating the relationship using appraisal values as substitutes for price expectations, the estimate is close to zero and not statistically significant. The difference in estimates from the two approaches is assessed by relaxing the identifying assumption that unobserved heterogeneity in housing units is time-invariant. This adjustment makes biases from omitted variables less tractable than those found in previous studies.

The difference in estimates is examined empirically in a step-wise manner. This is done for both time-invariant and time-varying heterogeneity in an effort to investigate their importance for identification. The results from this investigation suggests two new insights: first, time-varying unobserved heterogeneity plays an important role in biasing estimates when substituting with hedonic predictions of selling prices; second, the problem of omitted variable bias should be considered in relation to multiple explanatory variables suffering from differential measurement error, making the direction and magnitude of biases intractable.

In contrast to the existing literature, the findings suggest that sellers in the Oslo market do not exhibit nominal loss aversion in their choice of list prices. This result should be interpreted as being dependent on the institutional context in which sellers operate. First, sellers hire experts to provide appraisal values as tools for deciding on list prices, which may influence expectations through mechanisms such as anchoring. Second, buyers can observe the appraisal values, so even though sellers are loss averse, they may not act accordingly because this will provide signals about their types. Thus, the estimate accounts for and highlights the importance the institutional context. Future research of seller loss aversion should take the institutional context into consideration when collect more suitable data for estimating the effect, and this data should measure sellers' price expectations at the time of listing their housing units on the market.

References

Andersen, S., Badarinza, C., Liu, L., Marx, J., & Ramadorai, T. (2022). Reference Dependence in the Housing Market. *American Economic Review*, *112*(10), 3398–3440.

Anenberg, E. (2011). Loss aversion, equity constraints and seller behavior in the real estate market. *Regional Science and Urban Economics*, *41*(1), 67–76.

Anundsen, A. K., Nenov, P., Larsen, E. R., & Sommervoll, D. E. (2022). *Pricing and incentives in the housing market* (Housing Lab Working Paper No. 3). Oslo Metropolitan University.

Anundsen, A. K., & Røed Larsen, E. (2018). Testing for Micro-Efficiency in the Housing Market. *International Economic Review*, *59*(4), 2133–2162.

Beggs, A., & Graddy, K. (2009). Anchoring effects: Evidence from art auctions. *American Economic Review*, *99*(3), 1027–39.

Berkson, J. (1950). Are there two regressions? *Journal of the American Statistical Association*, *45*(250), 164–180.

Bokhari, S., & Geltner, D. (2011). Loss aversion and anchoring in commercial real estate pricing: Empirical evidence and price index implications. *Real Estate Economics*, *39*(4), 635–670.

Bracke, P., & Tenreyro, S. (2021). History Dependence in the Housing Market. *American Economic Journal: Macroeconomics*, *13*(2), 420–43.

Case, K. E., & Shiller, R. J. (1987). *Prices of single family homes since 1970: New indexes for four cities* (No. w2393). National Bureau of Economic Research.

Clapp, J. M., & Zhou, T. (2020). Controlling unobserved heterogeneity in repeat sales models: Application to anchoring to purchase price. *Available at SSRN 3358401*. https://doi.org/10.2139/ssrn.3358401

Einiö, M., Kaustia, M., & Puttonen, V. (2008). Price setting and the reluctance to realize losses in apartment markets. *Journal of Economic Psychology*, *29*(1), 19–34.

Genesove, D., & Mayer, C. (2001). Loss Aversion and Seller Behavior: Evidence from the Housing Market. *The Quarterly Journal of Economics*, *116*(4), 1233–1260.

Golub, G. H., & Van Loan, C. F. (2013). *Matrix computations* (2nd ed.). Johns Hopkins University Press.

Graddy, K., Loewenstein, L., Mei, J., Moses, M., & Pownall, R. A. (2022). Empirical evidence of anchoring and loss aversion from art auctions. *Journal of Cultural Economics*, 1–23.

Haber, G., Sampson, J., & Graubard, B. (2021). Bias due to Berkson error: issues when using predicted values in place of observed covariates. *Biostatistics*, *22*(4), 858–872.

Kahneman, D., & Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, *47*(2), 263.

Kahneman, D., & Tversky, A. (1984). Choices, values, and frames. *American Psychologist*, *39*(4), 341–350.

Lamorgese, A. R., & Pellegrino, D. (2022). Loss aversion in housing appraisal: Evidence from Italian homeowners. *Journal of Housing Economics*, *56*, 101826.

Norsk takst. (2022). *Etiske retningslinjer og regler for god takseringsskikk*. Retrieved November 6, 2022, from https://www.norsktakst.no/norsk/om-norsk-takst/regelverk/etiske-retningslinjer-og-regler-for-god-takseringsskikk/

Northcraft, G. B., & Neale, M. A. (1987). Experts, Amateurs, and Real Estate: An Anchoring-and- Adjustment Perspective on Property Pricing Decisions. *Organizational Behavior and Human Decision Processes*, *39*(1), 84–97.

Realtor.com. (2022). *Housing inventory: Median days on market in the united states* [From Federal Reserve Bank of St. Louis]. Retrieved November 6, 2022, from https://fred.stlouisfed.org/series/MEDDAYONMARUS

Strahilevitz, M. A., & Loewenstein, G. (1998). The effect of ownership history on the valuation of objects. *Journal of consumer research*, *25*(3), 276–289.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, *185*(4157), 1124–1131.

Zhou, T., Clapp, J. M., & Lu-Andrews, R. (2022). Examining omitted variable bias in anchoring premium estimates: Evidence based on assessed value. *Real Estate Economics*, *50*(3), 789–828.

A More on the data and hedonic regressions

A.1 On the main hedonic predictions

The hedonic predictions of selling prices are produced by a backward-looking procedure, reflecting how sellers are informed about the prior selling prices, as they are not able to look into the future. In each period (year-by-quarter), transactions from the earlier periods including the final period are used to fit a hedonic model. This model is used to predict selling prices for listings in that last period. Inclusion of the last period is intentional to estimate an implicit price for the time-effect in that last period, otherwise, predictions would tend to undervalue the units because aggregate prices have mostly increased during the sample period.[38]

The specification for the hedonic model has the log selling price as the dependent variable. When omitting subscripts, the model has the following specification:

$$\begin{aligned} log(P) = {} & \beta_0 + \beta_1 log(\text{size}) + \beta_2 log(\text{size})^2 + \beta_3 \mathbb{1}\{\text{apartment}\} \\ & + \beta_4 \mathbb{1}\{\text{apartment}\} \cdot log(\text{size}) + \beta_5 \mathbb{1}\{\text{apartment}\} \cdot log(\text{size})^2 \\ & + \beta_6 \mathbb{1}\{\text{built 1900-1949}\} + \beta_7 \mathbb{1}\{\text{built 1950-1979}\} + \beta_8 \mathbb{1}\{\text{built 1980-1999}\} \\ & + \beta_9 \mathbb{1}\{\text{built 2000-}\} + \beta_{10} \mathbb{1}\{\text{lot size} \geq 1000m^2\} \\ & + \beta_{11} \mathbb{1}\{\text{lot size} \geq 1000m^2\} \cdot \mathbb{1}\{\text{apartment}\} \\ & + \beta_{12} \mathbb{1}\{\text{non-co-op}\} + \beta_{13} \mathbb{1}\{\text{non-co-op}\} \cdot \mathbb{1}\{\text{apartment}\} \\ & + \sum_k \gamma_k \mathbb{1}\{\text{Zip code} = k\} + \sum_{t=2003Q2}^{T} \delta_t \mathbb{1}\{\text{Quarter sold} = t\} + e, \end{aligned} \tag{A.1}$$

which is estimated with $T =$ Q1 2005 as the first period to predict. Thus, the first hedonic model fit is the one that uses transactions for $t \leq$ Q1 2005 in the estimation. For the year-by-quarter dummies, Q1 2003 is the omitted period. To establish how well the model fits the sample, the model fit is evaluated using R^2 when predicting log selling prices for the sales happening in the same period as they are listed. Notably, the majority of the transactions occur within the same listing quarter. Figure A.1 shows R^2 for the model fits and the predictions when list and selling quarters are the same. Mechanically, the fitted R^2 curve is increasing when including more variables: when increasing the sample period, more time-dummies are included, and more observations may result in more zip code dummies.

[38] A work-around could be to estimate an AR-model for the index of each hedonic fit, and adjust predictions according to the AR-predictions of the consecutive period.

Figure A.1: R^2 for backward-looking hedonic estimates

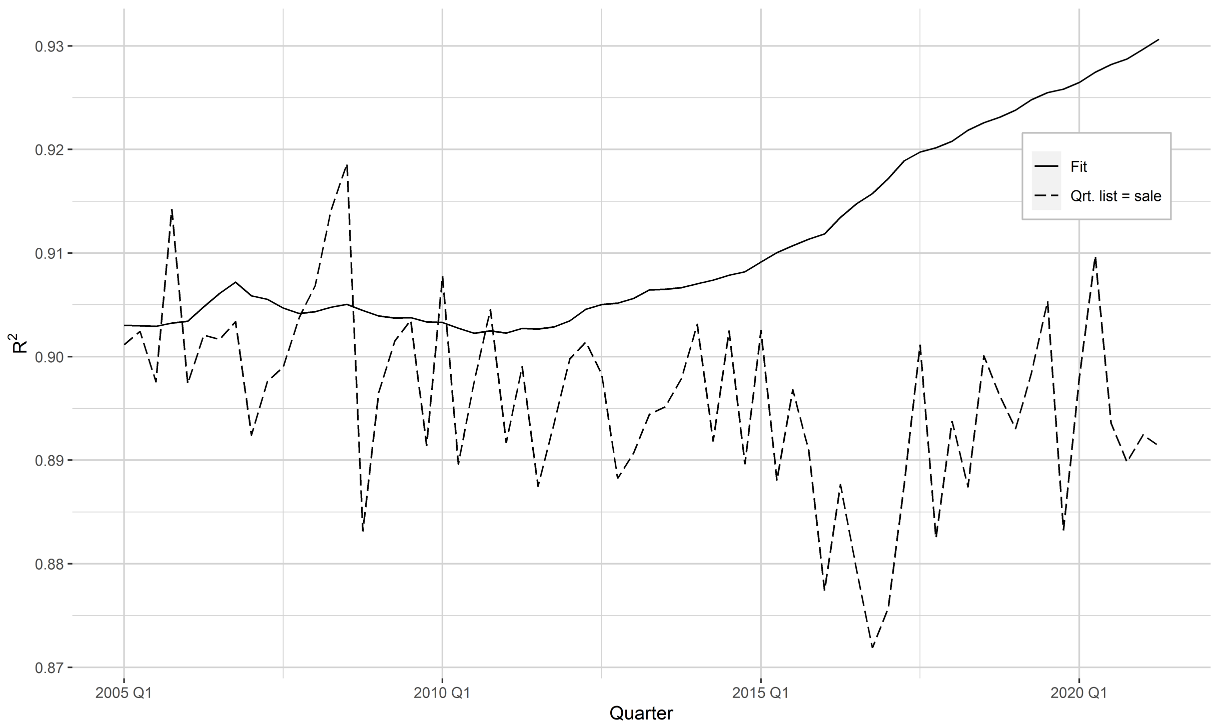

Notes: The figure presents the R^2 for the model fits and the predictions when the list and selling quarters are the same.

A.2 On the other predictions

The analysis utilizes other fitted values in addition to predictions of selling prices. All additional models used to produce fitted values use the same explanatory variables as equation (A.1), but they use a standard pooling OLS approach. Still, there are some deviations. First, the GM replication uses the specification with listing quarter dummies in place for selling quarter dummies, which replicates GM. The index used as a control variable in Table 3 is based on the dummy coefficients from estimating this model. Second, the index presented in Figure 1 and Figure D.3 uses a pooling approach. Third, the model generating fitted appraisal values uses the log appraisal value as the dependent variable and controls for listing-quarter dummies, because appraisal values are presented to sellers prior to listing. Fourth, the model that adds the appraisal value residuals uses the residuals from the hedonic appraisal values model as an explanatory variable. This model controls for selling year-quarter dummies and is fitted for listing year-quarter.

B Additional Derivations

B.1 Review of biases including the prospective gain

This appendix reviews model specification (3). The reduced form of the ideal specification is

$$\begin{aligned} L_{ist} = \alpha_0 + \alpha_1 \left(X_i\beta + \delta_t + v_i\right) + m_l \left(\delta_s - \delta_t + w_{is}\right)^+ \\ + m_g \left(\delta_s - \delta_t + w_{is}\right)^- + \epsilon_{it}. \end{aligned} \tag{B.1}$$

The feasible specification without adding the residual control is

$$\begin{aligned} L_{ist} = \alpha_0 + \alpha_1 \left(X_i\beta + \delta_t\right) + m_l \left(\delta_s - \delta_t + v_i + w_{is}\right)^+ \\ + m_g \left(\delta_s - \delta_t + v_i + w_{is}\right)^- + \eta_{it}, \end{aligned} \tag{B.2}$$

and the residual is

$$\begin{aligned} \eta_{it} &= \alpha_1 v_i + m_l(LOSS^* - LOSS) + m_g(GAIN^* - GAIN) + \epsilon_{it} \\ &= \alpha_1 v_i + m_l \left(\left(\delta_s - \delta_t + w_{is}\right)^+ - \left(\delta_s - \delta_t + v_i + w_{is}\right)^+ \right) \\ &\quad + m_g \left(\left(\delta_s - \delta_t + w_{is}\right)^- - \left(\delta_s - \delta_t + v_i + w_{is}\right)^- \right) + \epsilon_{it}. \end{aligned} \tag{B.3}$$

When assuming that unobserved quality does not change over time, $v_{is} = v_{it} = v_i$, the first part of the residual is contributing with a positive bias in the estimates, while the second part comes from the additional noise from the unobserved heterogeneity. Equation (B.2) shows that there should be a negative correlation between the price prediction and the prospective loss and gain due to δ_t. Furthermore, α_1 should be close to 1, potentially leading to a relatively strong positive bias, likely yielding upward biased estimates.

When including the residual from the hedonic prediction of the log selling price at the previous sale, and still assuming $v_{is} = v_{it} = v_i$, the reduced form becomes:

$$\begin{aligned} L_{ist} = \alpha_0 + \alpha_1 \left(X_i\beta + \delta_t\right) + \alpha_1 \left(v_i + w_{is}\right) \\ + m_l \left(\delta_s - \delta_t + v_i + w_{is}\right)^+ + m_g \left(\delta_s - \delta_t + v_i + w_{is}\right)^- + u_{it}, \end{aligned} \tag{B.4}$$

and the residual is

$$u_{it} = -\alpha_1 w_{is} + m_l \left(LOSS^*_{ist} - LOSS_{ist}\right) + m_g \left(GAIN^*_{ist} - GAIN_{ist}\right) + \epsilon_{it}. \tag{B.5}$$

The residual again shows two sources of potential bias in m_l and m_g. When including the residual control variable, $(v_i + w_{is})$, u_{it} includes $-\alpha_1 w_{is}$ which should contribute with a downward bias. The second part of u_{it} remains unchanged, while noting that w_{is} is included in the feasible loss and gain and the residual control. The residual control should correlate positively with loss and gain, leading to possible direct effects on m_l and m_g, but also indirect effects from the $-\alpha_1 w_{is}$ in u_{it} through the residual control. In total, we may get downward biased estimates.

Finally, if not invoking constant unobserved heterogeneity, the ideal model becomes as presented in equation (5), and the feasible model becomes

$$L_{ist} = \alpha_0 + \alpha_1 \left(X_i\beta + \delta_t\right) + m_l \left(\delta_s - \delta_t + v_{is} + w_{is}\right)^+ + m_g \left(\delta_s - \delta_t + v_{is} + w_{is}\right)^- + \eta_{it}, \quad \text{(B.6)}$$

and the residual is

$$\eta_{it} = \alpha_1 v_{it} + m_l(LOSS^* - LOSS) + m_g(GAIN^* - GAIN) + \epsilon_{it}. \quad \text{(B.7)}$$

This should again yield upward biased estimates. To bias the estimates downward, the lagged residual is again added making the reduced-form of the residual as follows:

$$u_{it} = \alpha_1(v_{it} - v_{is}) - \alpha_1 w_{is} + m_l(LOSS^* - LOSS) + m_g(GAIN^* - GAIN) + \epsilon_{it}, \quad \text{(B.8)}$$

so that the change in unobserved heterogeneity is part of the residual. The estimates of interest are likely downward biased from the terms in the residual, while correlations between explanatory variables could lead to upward biases in the estimates. With time-varying unobserved heterogeneity, the downward biasing effect may be smaller than with time-invariant unobserved heterogeneity, because the unobserved heterogeneity at time t does not disappear from the residual (B.8). Therefore, it is not clear whether the lower bound actually is a lower bound.

B.2 Differential measurement error in loss aversion models

B.2.1 Baseline

In what follows, the general form of the Berkson bias is given as presented by Haber et al. (2021). Leaving out subscripts $\{i, s, t\}$, let the causal relationship be

$$L = \lambda_0 + \lambda_1\mu + \lambda_l LOSS + \lambda_\pi \pi + \epsilon, \quad \text{(B.9)}$$

while the estimated model is

$$L = \alpha_0 + \alpha_1\mu + m_l LOSS + \eta, \quad \text{(B.10)}$$

in which $\hat{P}$ is substituted for μ. The Berkson model is $\mu = \hat{P} + e$, with e as the Berkson residual. Note that the causal relationship includes the observed prospective loss and not the ideal version. This is done deliberately to include the mismeasured LOSS that has v_{is} in it.[39] By result 4.1 in Haber et al. (2021), the bias in $\hat{m}$ from estimating B.10 with $\hat{P}$ instead of μ is

$$\begin{aligned}
&E\left(\hat{m}_l(\hat{P}) - \hat{m}_l(\mu)\right) = A^{-1}B \times \left(-(\sigma_{\mu,l} - \sigma_{e,l})\sigma_l^{-2}\right) + \alpha_1\sigma_{e,l}\sigma_l^{-2} \\
&A = \sigma_\mu^2 - \sigma_e^2 - (\sigma_{\mu,l} - \sigma_{e,l})^2\sigma_l^{-2} \\
&B = (\alpha_1 - \lambda_1)\sigma_e^2 + (m_l - \lambda_l)\sigma_{e,l} - \lambda_\pi\sigma_{\pi,e} - \alpha_1\sigma_{e,l}\sigma_l^{-2}(\sigma_{\mu,l} - \sigma_{e,l}).
\end{aligned} \tag{B.11}$$

The notation $\hat{m}_l(\hat{P})$ means that equation (B.10) is estimated with $\hat{P}$, while $\hat{m}_l(\mu)$ denotes the estimate if the estimated relationship used μ. Also, the subscript l denotes the LOSS, so that $Var(LOSS) = \sigma_l^2$, and so on. These relations are useful:

$$\begin{aligned}
\alpha_1 - \lambda_1 &= \lambda_\pi \frac{\sigma_{\pi,\mu|l}}{\sigma^2_{\mu|l}} = \lambda_\pi \frac{\sigma_{\pi,\mu}\sigma_l^2 - \sigma_{\pi,l}\sigma_{\mu,l}}{\sigma_\mu^2\sigma_l^2 - \sigma_{\mu,l}^2} \\
m_l - \lambda_l &= \lambda_\pi \frac{\sigma_{\pi,l|\mu}}{\sigma^2_{l|\mu}} = \lambda_\pi \frac{\sigma_{\pi,l}\sigma_\mu^2 - \sigma_{\pi,\mu}\sigma_{\mu,l}}{\sigma_\mu^2\sigma_l^2 - \sigma_{\mu,l}^2}.
\end{aligned}$$

These reflect that the estimates from estimating equation (B.10) with μ only differs if there is no omitted variable bias. Assume that $\sigma_{\pi,\mu} = \sigma_{\pi,l} = 0$, which gives $\alpha_1 = \lambda_1$ and $m_l - \lambda_l$. Using this in equation (B.11), the bias becomes

$$\begin{aligned}
E\left(\hat{m}_l(\hat{P}) - \hat{m}_l(\mu)\right) = \lambda_\pi &\frac{\sigma_{\pi,e}(\sigma_{\mu,l} - \sigma_{e,l})}{(\sigma_\mu^2 - \sigma_e^2)\sigma_l^2 - (\sigma_{\mu,l} - \sigma_{e,l})^2} \\
&+ \alpha_1 \frac{\sigma_{e,l}}{\sigma_l^2} \frac{(\sigma_\mu^2 - \sigma_e^2)\sigma_l^2}{(\sigma_\mu^2 - \sigma_e^2)\sigma_l^2 - (\sigma_{\mu,l} - \sigma_{e,l})^2}.
\end{aligned} \tag{B.12}$$

The first part of the expression is the bias from the correlation between the Berkson error and the unobserved variable, and the second part is from the correlation between the Berkson error and prospective loss.

B.2.2 Adding hedonic residual

The expression in equation (B.11) is the single variable reduction of a more general expression for the bias, meaning, there is only one variable in addition to the mismeasured variable included in the estimated model (B.10). If the hedonic residual $\epsilon_{is} = v_{is} + w_{is}$ is included, here denoted by ε, from the hedonic selling price prediction of the previous sale, this adds one variable consisting of two elements to the model. The bias expression of estimating the model, with notation $\boldsymbol{z'} = (l, \varepsilon)$ and $\beta_z' = (m_l, \beta_\varepsilon)$, while keeping to one omitted variable and maintaining

[39] v_{is} makes the Berkson error *differential*.

the assumptions for the expression (B.11), and denoting the covariance matrix as Σ, becomes

$$\begin{aligned} E\left(\hat{\beta}_z(\hat{P}) - \hat{\beta}_z(\mu)\right) = &\left[\sigma_\mu^2 - \sigma_e^2 - (\Sigma_{\mu,z} - \Sigma_{e,z})\Sigma_{z,z}^{-1}(\Sigma_{z,\mu} - \Sigma_{z,e})\right]^{-1} \\ &\left[-\lambda_\pi \sigma_{\pi,e} - \alpha_1 \Sigma_{e,z}\Sigma_{z,z}^{-1}(\Sigma_{z,\mu} - \Sigma_{z,e})\right] \\ &\left(-(\Sigma_{\mu,z} - \Sigma_{e,z})\Sigma_{z,z}^{-1}\right) + \alpha_1 \Sigma_{e,z}\Sigma_{z,z}^{-1}. \end{aligned}$$

The bias expression for $\hat{m}_l$ is found by solving the matrix form of the expression above, while noting that $C \geq 0$:

$$\begin{aligned} &E\left(\hat{m}_l(\hat{P}) - \hat{m}_l(\mu)\right) = C^{-1}\left[\alpha_1(\sigma_{e,l}\sigma_\varepsilon^2 - \sigma_{e,\varepsilon}\sigma_{l,\varepsilon}) + DE^{-1}\left((\sigma_{\mu,l} - \sigma_{e,l})\sigma_\varepsilon^2 - (\sigma_{\mu,\varepsilon} - \sigma_{e,\varepsilon})\sigma_{l,\varepsilon}\right)\right] \\ &C = \sigma_l^2\sigma_\varepsilon^2 - \sigma_{l,\varepsilon}^2 = \sigma_{l,\varepsilon}^2\left(\frac{1}{\rho_{l,\varepsilon}^2} - 1\right) \quad \text{if} \quad \rho_{l,\varepsilon} \neq 0 \\ &D = \alpha_1\left(\sigma_{e,l}\sigma_\varepsilon^2(\sigma_{\mu,l} - \sigma_{e,l}) + \sigma_{e,\varepsilon}\sigma_l^2(\sigma_{\mu,\varepsilon} - \sigma_{e,\varepsilon}) + \sigma_{l,\varepsilon}(2\sigma_{e,l}\sigma_{e,\varepsilon} - \sigma_{e,\varepsilon}\sigma_{\mu,l} - \sigma_{e,l}\sigma_{\mu,\varepsilon})\right) \\ &\qquad + \lambda_\pi\sigma_{\pi,e}(\sigma_l^2\sigma_\varepsilon^2 - \sigma_{l,\varepsilon}^2) \\ &E = (\sigma_\mu^2 - \sigma_e^2)(\sigma_l^2\sigma_\varepsilon^2 - \sigma_{l,\varepsilon}^2) - (\sigma_{\mu,l} - \sigma_{e,l})^2\sigma_\varepsilon^2 - (\sigma_{\mu,\varepsilon} - \sigma_{e,\varepsilon})^2\sigma_l^2 + 2(\sigma_{\mu,l} - \sigma_{e,l})(\sigma_{\mu,\varepsilon} - \sigma_{e,\varepsilon})\sigma_{l,\varepsilon}. \end{aligned}$$

B.2.3 Measurement error in prospective loss

When considering a case where μ is perfectly measured, but $LOSS^*$ is measured as $LOSS$, the causal relationship is $L = \lambda_0 + \lambda_1\mu + \lambda_l LOSS^* + \lambda_\pi\pi + \epsilon$. The estimated relationship is $L = \alpha_0 + \alpha_1\mu + m_l LOSS^* + \eta$, with $LOSS$ instead of $LOSS^*$. Denote $LOSS^* - LOSS = e_l$, then the bias expression is

$$E\left(\hat{m}_l(\hat{P}) - \hat{m}_l(\mu)\right) = -\frac{\lambda_\pi\sigma_\mu^2\sigma_{\pi,e_l} + m_l\sigma_{\mu,e_l}(\sigma_{l,\mu} - \sigma_{\mu,e_l})}{(\sigma_l^2 - \sigma_{e_l}^2)\sigma_\mu - (\sigma_{l,\mu} - \sigma_{\mu,e_l})^2},$$

in which the covariance $\sigma_{\mu,e_l} \neq 0$ because

$$e_l = \Delta L = (\delta_s - \delta_t + v_{is} - v_{it} + w_{is})^+ - (\delta_s - \delta_t + v_{is} + w_{is})^+.$$

This difference comes from the missing v_{it}, a term that also is included in μ_{it}.

C More on the appraisal values

C.1 List prices and expectation substitutes

A key concern in the paper relates to the choice of price expectation substitutes, namely how well the two different substitutes explain the list price and which to possibly prefer. Another concern is whether there is bunching, particularly between the list price and appraisal value. As presented in Table C.1, both hedonic predictions and appraisal values have high explanatory power on the list price.[40] Even though predictions are explaining log selling prices well, appraisal values are even better. When comparing with results from the loss aversion models, the additional variation explained by the prospective terms and control variables are limited. Moreover, when comparing different spreads between list prices (selling prices) and the substitutes in Figures C.1 and C.2, hedonic predictions show approximately normally distributed spreads, while the appraisal value spreads have strong bunching at zero. There are (almost) no positive spreads for the list price to appraisal values spread, which together with the strong zero bunching suggests that appraisal values influence the choice list price.[41]

The last column in Table C.1 shows what happens when adding both of the substitutes (appraisal values and hedonic predictions) to explain list prices, resulting in a large coefficient estimate on the appraisal values and a small on the hedonic predictions. Together with the higher explanatory power of appraisal values, this is evidence in favor of using appraisal values as substitutes for the *main list price determinant*. However, in the last column, which includes both substitutes, the case is less clear-cut than it may seem: there is a strong correlation between the two alternatives because both have common components which are directly relevant for list prices. To address this, the two substitutes are orthogonalized using the modified Gram-Schmidt procedure (Golub & Van Loan, 2013, pp. 254–255). In this procedure, the order of which the orthogonalization is performed is important for outcomes, making it necessary to repeat the procedure twice, treating each substitute as the primary substitute.[42] This means that the final orthogonal variables are different depending on the order of which they appear in the matrix that will be orthogonalized. Also, note that all variables in the matrix are normalized, so that the coefficients can no longer be directly interpreted.

Table C.3 presents the results from orthogonalizing appraisal values and hedonic predictions. An ad-hoc test of their relevance for the list price is to simply compare their coefficient sizes, meaning the primary variable relative to the secondary variable. This allows to test whether the coefficient estimate of the first column in the orthogonalized matrix is larger than the second

[40] Table C.2 provides results from investigating the explanatory power of the different substitutes on the selling price.

[41] The spreads in Figure C.1 are trimmed on the 1st and 99th percentiles in order to get a cleaner visualization. There are 26 observations in the sample with the list price above the appraisal value.

[42] The primary variable is the first column of the matrix containing columns that are orthogonalized.

column.

Because the a priori strongest contender for explaining the list price is the appraisal value, results from performing the procedure with appraisal values in the first column in the matrix that is orthogonalized is presented first. Here, the hedonic predictions are orthogonalized on the appraisal values. The coefficient estimate on appraisal value is about 100 times larger than the coefficient estimate on the predictions, but this should be interpreted jointly with the second half of the table, which does the procedure with hedonic predictions as the primary variable. Here, the first column in the orthogonalized matrix contains the predictions, which yields a coefficient on the predictions that is 3.3 times larger than the coefficient on appraisal values. Although a smaller fraction, the p-value of testing whether this relative size is less than one suggests rejecting the null hypothesis. Thus, with this ad-hoc test of choosing between substitutes, the relative sizes of the coefficients in the two cases suggest that the appraisal values may be preferred to hedonic predictions in explaining list prices. Together with results in Table C.1, appraisal values should be preferred over hedonic predictions as the *main list price determinant*.

Table C.1: Log list prices on (predicted) log prices

	(1)	(2)	(3)	(4)	(5)
$\hat{P}_{it}$	1.003***				0.048***
	(0.012)				(0.007)
P_{it}^{AV}		1.001***			0.957***
		(0.002)			(0.007)
$\hat{P}_{it}^{AV}$			0.999***		
			(0.010)		
$\hat{P}_{it}\|_{e_{it}^{AV}}$				1.020***	
				(0.007)	
Constant	−0.091	−0.029	0.011	−0.344***	−0.086***
	(0.171)	(0.028)	(0.144)	(0.109)	(0.030)
N	15,506	15,506	15,506	15,506	15,506
Adj. R sq.	0.923	0.996	0.926	0.990	0.996

Notes: The table presents results regressing log(List price) on backward-looking hedonic predictions ($\hat{P}_{it}$), log(Appraisal values) (P_{it}^{AV}), fitted log(Appraisal values) ($\hat{P}_{it}^{AV}$), and cross-sectional hedonic predictions with appraisal value residuals ($\hat{P}_{it}|_{e_{it}^{AV}}$). All estimations use the same observations from the appraisal sample. Standard errors are clustered on list year and 3-digit zip codes. Significance: * p<0.1, ** p<0.05, *** p<0.01.

Table C.2: Log selling prices on (predicted) log prices

	(1)	(2)	(3)	(4)	(5)
$\hat{P}_{it}$	0.983***				0.156***
	(0.012)				(0.031)
P_{it}^{AV}		0.971***			0.827***
		(0.010)			(0.024)
$\hat{P}_{it}^{AV}$			0.971***		
			(0.014)		
$\hat{P}_{it}\|_{e_{it}^{AV}}$				0.997***	
				(0.007)	
Constant	0.273	0.464***	0.487**	0.049	0.277*
	(0.173)	(0.147)	(0.213)	(0.099)	(0.163)
N	15,506	15,506	15,506	15,506	15,506
Adj. R sq.	0.909	0.963	0.897	0.972	0.965

Notes: The table presents results regressing log(Selling price) on backward-looking hedonic predictions ($\hat{P}_{it}$), log(Appraisal values) (P_{it}^{AV}), fitted log(Appraisal values) ($\hat{P}_{it}^{AV}$), and cross-sectional hedonic predictions with appraisal value residuals ($\hat{P}_{it}|_{e_{it}^{AV}}$). All estimations use the same observations from the appraisal sample. Standard errors are clustered on selling year and 3-digit zip codes. Significance: * p<0.1, ** p<0.05, *** p<0.01.

Figure C.1: List price spreads

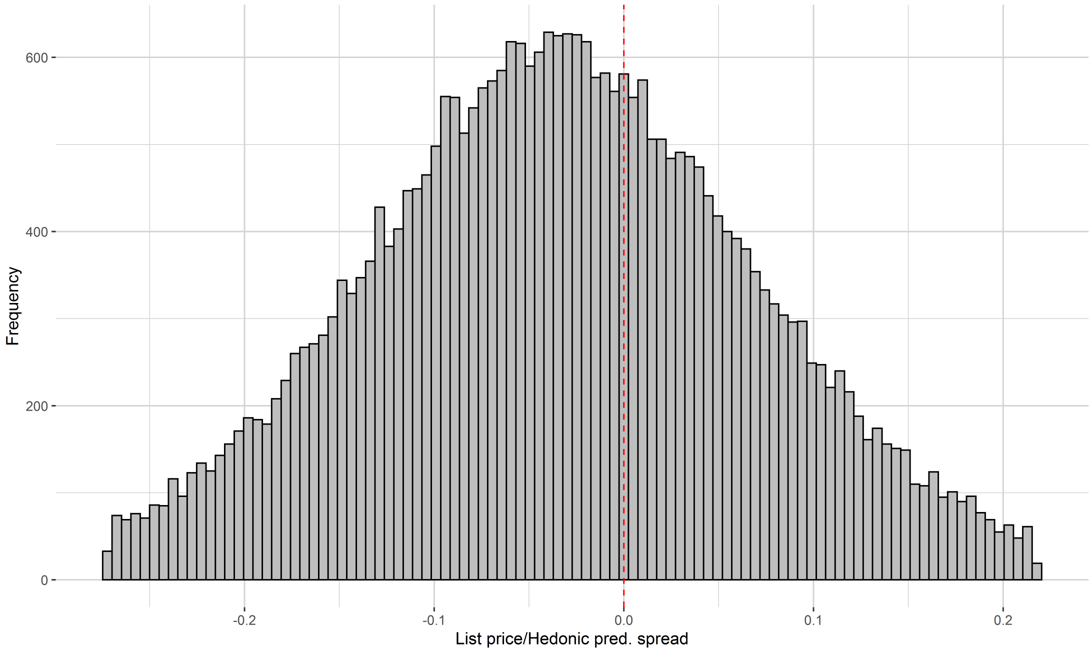

(A) List price/hedonic spread

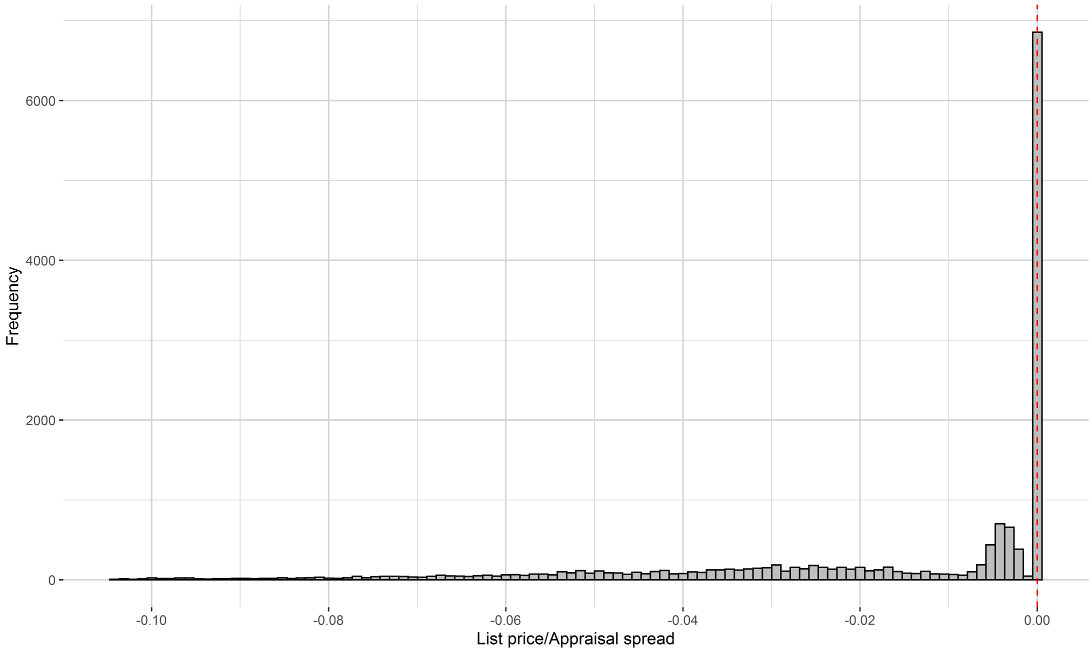

(B) List price/appraisal value spread

Notes: The figure presents the spreads, found by the difference in log terms: log list price minus predicted log selling price, and log list price minus log appraisal value. Panel A is the list price/hedonic prediction spread from the hedonic sample, and panel B the list price/appraisal value spread from the appraisal sample. For a clean visualization, the spreads are trimmed on the 1st and 99th percentiles.

Figure C.2: Selling price spreads

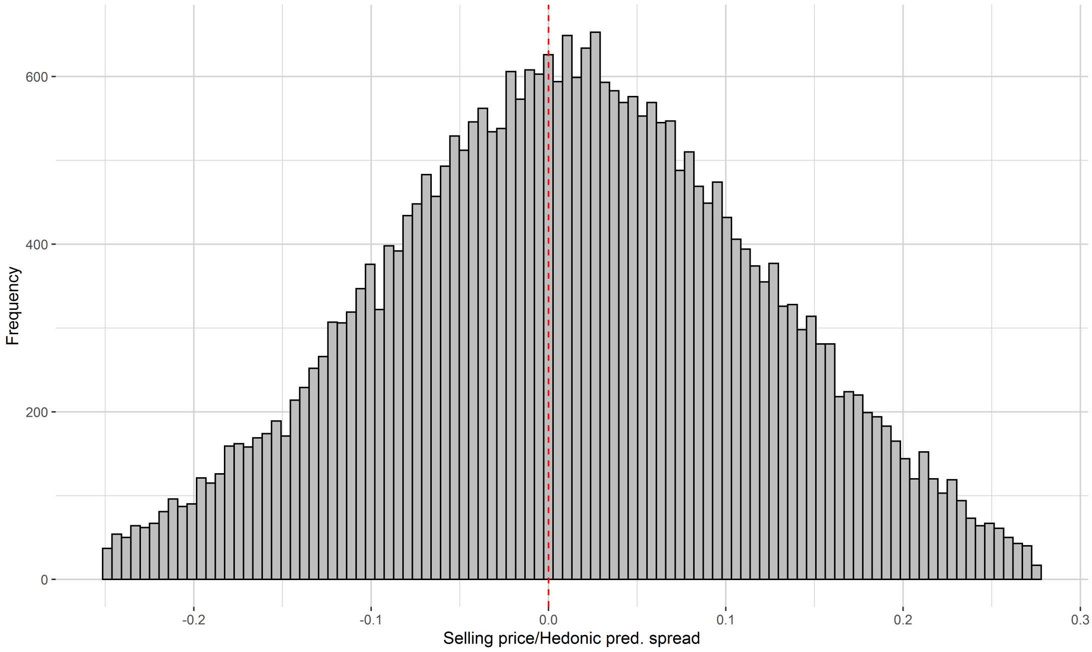

(A) Selling price/hedonic spread

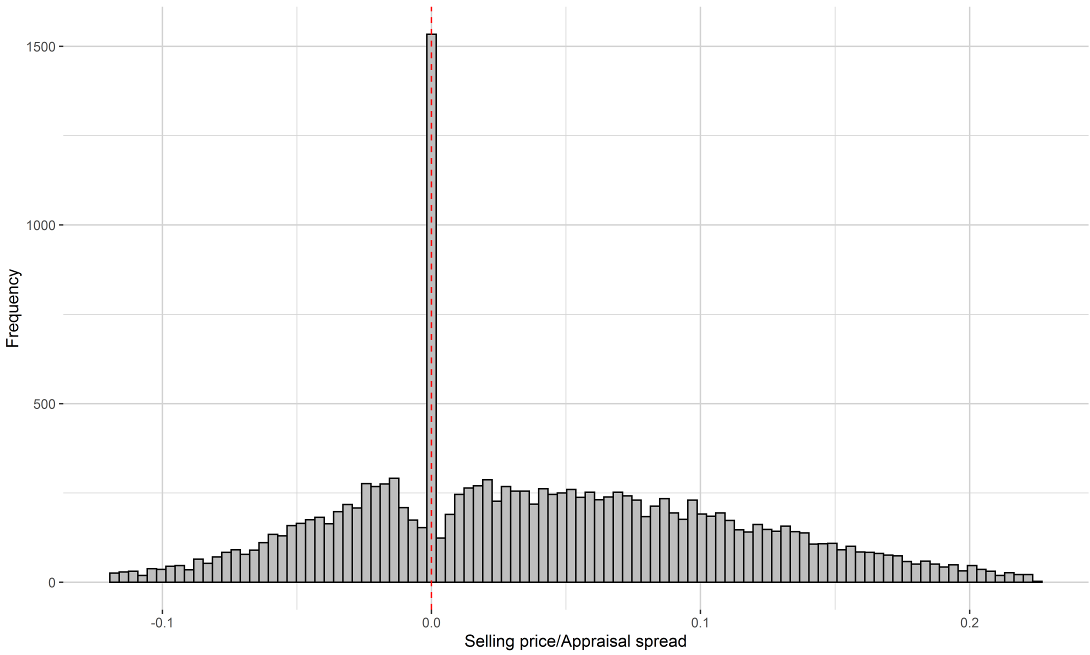

(B) Selling price/appraisal value spread

Notes: The figure presents the spreads, found by the difference in log terms: log selling price minus predicted log selling price, and log selling price minus log appraisal value. Panel A is the selling price/hedonic prediction spread from the hedonic sample, and panel B the selling price/appraisal value spread from the appraisal sample. For a clean visualization the spreads are trimmed on the 1st and 99th percentiles.

Table C.3: Orthogonalizing appraisal values and hedonic predictions

	1st col.: AV			1st col.: Hed.pred.		
	(1)	(2)	(3)	(4)	(5)	(6)
log(Appraisal value)	0.396***		0.396***	0.114***		0.114***
	(0.001)		(0.001)	(0.016)		(0.001)
Predicted log(Selling price)		0.004	0.004***		0.379***	0.379***
		(0.015)	(0.001)		(0.005)	(0.001)
Constant	14.856***	14.856***	14.856***	14.856***	14.856***	14.856***
	(0.002)	(0.039)	(0.002)	(0.044)	(0.008)	(0.002)
N	15,884	15,884	15,884	15,884	15,884	15,884
Adj. R sq.	0.996	0.000	0.996	0.082	0.914	0.996
1st col./2nd col.			99.910			3.333
p-value of H_0: 1st col.$\leq$2nd col.			<0.001			<0.001

Notes: The table presents results from orthogonalizing log(Appraisal value) and hedonic predictions of log(Selling price), using the appraisal sample (with sample period Q1 2005 to Q2 2016). The dependent variable is log(List price). The first three columns orthogonalize with the appraisal value as the first column in the orthogonalization procedure, and the last three columns use the hedonic predictions as the first column. Standard errors are clustered on list year and 3-digit zip codes. Significance: * p<0.1, ** p<0.05, *** p<0.01.

C.2 Appraisal value residuals and unobserved heterogeneity

Recall the reduced form of the appraisal value from Subsection 3.3, $P_{it}^{AV} = X_i\beta + \delta_t + \tilde{v}_{it}$, with $\tilde{v}_{it} = P_{it}^{AV} - \hat{P}_{it}^{AV}$. The appraisal value residual, $\tilde{v}_{it}$, is the unobserved heterogeneity in the appraisal value, consisting of both time-invariant and time-varying heterogeneity, such as proximity to amenities and time-varying quality. The appraisal value should be a more accurate measure of the true expectation than a hedonic prediction of selling price, because surveyors observe more information than what is included in the sample. Formally, $E|\psi_{it}| \leq E|v_{it}|$, in which $\psi_{it} = \mu_{it} - P_{it}^{AV}$ and $v_{it} = \mu_{it} - \hat{P}_{it}$.

The results in Table 8 suggest that when changing the main list price determinant in the baseline hedonic-based specification, the loss aversion effect disappears. This accounts for more of the unobserved heterogeneity in the predictions by using the full appraisal values, and possibly making the constructed loss and gain less endogenous. Following the notion of utilizing the additional information nested in appraisal values, represented by the appraisal value residuals, these are possible to split up into a time-invariant and time-varying part. Doing so allows for testing which part biasing the loss aversion estimate the most.

Specifically, $\tilde{v}_{it}$ is separated into two parts: the explained part and the unexplained part.[43] Let $\tilde{v}_{it} \equiv \tilde{v}_{it}^{ex} + \tilde{v}_{it}^{unex}$, with $\tilde{v}_{it}^{ex}$ being the explained part and $\tilde{v}_{it}^{unex}$ being the unexplained part of the residual. The appraisal residual model is

$$\begin{aligned}\tilde{v}_{it} = \beta_0 &+ \beta_1\tilde{v}_{is} + \beta_2 MAH_{ilt} + \beta_3 MAH_{ils} + \beta_4 log(\text{Holding time}_{it}) + \beta_5 log(\frac{\phi_{l,t}}{\phi_{l,t-4}}) \\ &+ \beta_6\tilde{v}_{is} \times MAH_{ils} + \beta_7\tilde{v}_{is} \times log(\text{Holding time}_{it}) + \beta_8 MAH_{ilt} \times log(\frac{\phi_{l,t}}{\phi_{l,t-4}}) \\ &+ \beta_9 MAH_{ilt} \times log(\text{Holding time}_{it}) + \tilde{e}_{it},\end{aligned} \tag{C.1}$$

which is used to obtain the fitted value $\hat{\tilde{v}}_{it} = \tilde{v}_{it}^{ex}$ and the residual $\hat{\tilde{e}}_{it} = \tilde{v}_{it}^{unex}$. Here, l denotes the location, so that MAH_{ilt} is the within list year and 3-digit zip code Mahalanobis distance compared to the respective means, calculated using the size, age and location (4-digit zip code). MAH_{ils} does the same for the year of the previous purchase (s). $\phi_{l,t}$ is the local repeat sales index, at the 3-digit zip code level, created using the Case-Shiller weighted repeated sales method (Case & Shiller, 1987). The repeat sales index is estimated using the *hedonic sample* restricted to the period Q1 2008-Q4 2020 due to missing observations in some 3-digit zip codes.[44] $\phi_{l,t-4}$ denotes the local index lagged four periods (quarters), making $log(\phi_{l,t}/\phi_{l,t-4})$ being the

[43] An alternative to this procedure is controlling for the $\tilde{v}_{it} - \tilde{v}_{is}$, meaning, the difference between the appraisal residuals between listing and previous sale. However, the presented approach also accounts for some potential noise as well.

[44] 3-digit zip codes with less than 500 observations are initially dropped from the full sample (listing in the period Q1 2005-Q2 2021). Observations with previous sales outside the sample period (Q1 2008-Q4 2020) are also dropped. Then observations before Q1 2008 and after Q4 2020 are dropped. This makes the sample size reduce from 32,044 to 22,194 repeated observations, which is the sample used to estimate the index.

appreciation in local selling prices compared to the same period the year before listing. Holding time is the usual holding time of the sellers measured in weeks.

Model (C.1) is supposed to capture why surveyors miss in their appraisal values. First, the residual from the previous sale should capture to what degree the same unobserved heterogeneity is still present in the non-lagged residual. Meaning, the autocorrelation between these should capture the time-invariant part. The model includes the local within list year Mahalanobis distance to capture heterogeneity in observables. The holding time is included as a proxy for time-varying heterogeneity, and the local appreciation should capture some of the noise in appraisal values attributed to changes in the seasonal price levels. Interactions with these variables are also included. The result from estimating the relationship is presented in the first column in Table C.5, showing that about 37.5 percent of the total variation in the appraisal residual is explained.

In Table C.4, the variables $\{\tilde{v}_{it}, \tilde{v}_{it}^{ex}, \tilde{v}_{it}^{unex}\}$ are used to control for more of the unobserved heterogeneity when estimating the baseline hedonic-based list price model. Due to the small sample resulting from using the local repeat sales index, the first column provides estimates of the baseline hedonic-based list price model for this subsample. When adding the full appraisal residual, the estimated loss aversion effect is reduced from 0.501 to 0.088.[45] Controlling for the explained part of the residual, consisting mostly of time-invariant unobserved heterogeneity and some additional noise, the estimate is almost at the same level of magnitude as the baseline estimate in the first column. Thus, the constant unobserved heterogeneity in appraisal values does not capture the unobserved heterogeneity in the predictions. However, adding the unexplained part makes the estimate drop to 0.172. Controlling for the unexplained residual implies controlling for mostly time-varying unobserved heterogeneity. Thus, this shows how the time-varying unobserved heterogeneity biases the estimate.

Furthermore, the results indicate that the time-varying unobserved heterogeneity in appraisal values partially captures the time-varying unobserved heterogeneity in hedonic predictions. Otherwise, this control variable should not have affected the estimate. Finally, the two last columns use the two residual parts as instruments with the full residual as the endogenous variable. Although these instruments are not valid, the exercise allows for using the variation from the two parts instead of the full residual. The same pattern emerges: when instrumenting with the unexplained part of the residual, the loss aversion estimate becomes smaller than when instrumenting with the explained part. Note that the prospective loss, the prospective gain and the predicted price are all considered exogenous and added as controls in the first stage, giving relatively strong explanatory power especially in the last column, making the loss aversion estimates differ from the previous columns. The first stage results are presented in the

[45] Investigating the potential bias as in the informal review in Appendix B.1, adding $\tilde{v}_{it}$ in equation (B.6), yields the residual $u_{it} = \alpha_1 \psi_{it} + m_l(LOSS^* - LOSS) + m_g(GAIN^* - GAIN) + \epsilon_{it}$. Thus, this leaves a much smaller biasing component in ψ_{it}, which should still bias the estimates upwards.

two last columns in Table C.5.

Table C.4: Accounting for unobserved variation in appraisal values

	OLS				IV	
	(1)	(2)	(3)	(4)	(5)	(6)
LOSS	0.951***	0.201***	0.784***	0.566***	0.361***	0.168***
	(0.034)	(0.019)	(0.033)	(0.031)	(0.026)	(0.020)
GAIN	0.450***	0.113***	0.313***	0.394***	0.185***	0.098***
	(0.010)	(0.005)	(0.012)	(0.006)	(0.009)	(0.005)
Predicted price	1.025***	1.013***	1.021***	1.023***	1.016***	1.013***
	(0.004)	(0.001)	(0.004)	(0.002)	(0.002)	(0.001)
Resid. (AV) ($\tilde{v}_{it}$)		0.858***			0.675***	0.895***
		(0.006)			(0.020)	(0.007)
Resid. (AV), explained ($\tilde{v}_{it}^{ex}$)			0.424***			
			(0.029)			
Resid. (AV), unexplained ($\tilde{v}_{it}^{unex}$)				0.798***		
				(0.009)		
Substitute for μ_{it}	$\hat{P}_{it}$	$\hat{P}_{it}$	$\hat{P}_{it}$	$\hat{P}_{it}$	$\hat{P}_{it}$	$\hat{P}_{it}$
Instrument					$\tilde{v}_{it}^{ex}$	$\tilde{v}_{it}^{unex}$
N	5,956	5,956	5,956	5,956	5,956	5,956
Adj. R sq.	0.946	0.989	0.949	0.981	0.987	0.989
LOSS>0 (% of N)	7.908	7.908	7.908	7.908	7.908	7.908
F-stat.(LOSS=GAIN)	171.157	19.071	178.100	27.847	57.330	11.752
p-value(LOSS=GAIN)	<0.001	<0.001	<0.001	<0.001	<0.001	0.001
LOSS-GAIN	0.501	0.088	0.472	0.172	0.176	0.070

Notes: The table presents results from utilizing the residual from the hedonic estimation of appraisal values, for the sample period Q1 2009 to Q2 2016. The dependent variable is log(List price), and hedonic predictions are used as substitutes for the price expectation. The first four columns are OLS estimates, with the first being the baseline estimation and the second adds the full residual as a control variable. The last two columns are IV estimates. The last four columns does the reverse. The log of holding time in weeks, DTV, and the constant term are omitted from the table. Standard errors are heteroskedasticity-robust (White) due to few cluster-groups. Significance: * p<0.1, ** p<0.05, *** p<0.01.

Table C.5: Separating the appraisal value residual

Initial stage	(1)	First stage	(2)	(3)
Resid (AV), prev. ($\tilde{v}_{is}$)	0.233**	LOSS	0.627***	0.444***
	(0.100)		(0.034)	(0.029)
MAH	-0.004	GAIN	0.190***	0.330***
	(0.004)		(0.012)	(0.005)
MAH, prev.	-0.003***	DTV	0.043***	0.006**
	(0.001)		(0.006)	(0.003)
log(Holding time)	-0.030***	Predicted price	0.007**	0.011***
	(0.004)		(0.004)	(0.002)
1 yr local appr.	0.081***	log(Holding time)	0.025***	0.029***
	(0.026)		(0.003)	(0.001)
Resid (AV), prev. ($\tilde{v}_{is}$) x MAH, prev.	0.009***	Resid (AV), ex. ($\tilde{v}_{it}^{ex}$)	0.629***	
	(0.003)		(0.029)	
Resid (AV), prev. ($\tilde{v}_{is}$) x log(Holding time)	0.053***	Resid (AV), unex. ($\tilde{v}_{it}^{unex}$)		0.892***
	(0.020)			(0.009)
MAH x 1 yr local appr.	-0.010*	Constant	-0.194***	-0.211***
	(0.006)		(0.052)	(0.027)
MAH x log(Holding time)	0.001			
	(0.001)			
Constant	0.171***			
	(0.018)			
N	5,956		5,956	5,956
Adj. R. sq.	0.375		0.435	0.835
F-stat. weak instr.			455.621	10818.117
p-value weak instr.			<0.001	<0.001

Notes: The table presents results with the residual from the hedonic estimation of appraisal values as dependent variable, for the sample period Q1 2009 Q1 to Q2 2016. The first column is the initial stage from which $\tilde{v}_{it}^{ex}$ and $\tilde{v}_{it}^{unex}$ are obtained, while the last two columns are the first stages for the instrumental variables results in Table C.4. Hedonic predictions are used as substitutes for the price expectations. Standard errors are heteroskedasticity-robust (White) due to few cluster-groups. Significance: * p<0.1, ** p<0.05, *** p<0.01.

D Additional Tables and Figures

Figure D.1: Price index and median time-on-market for Oslo

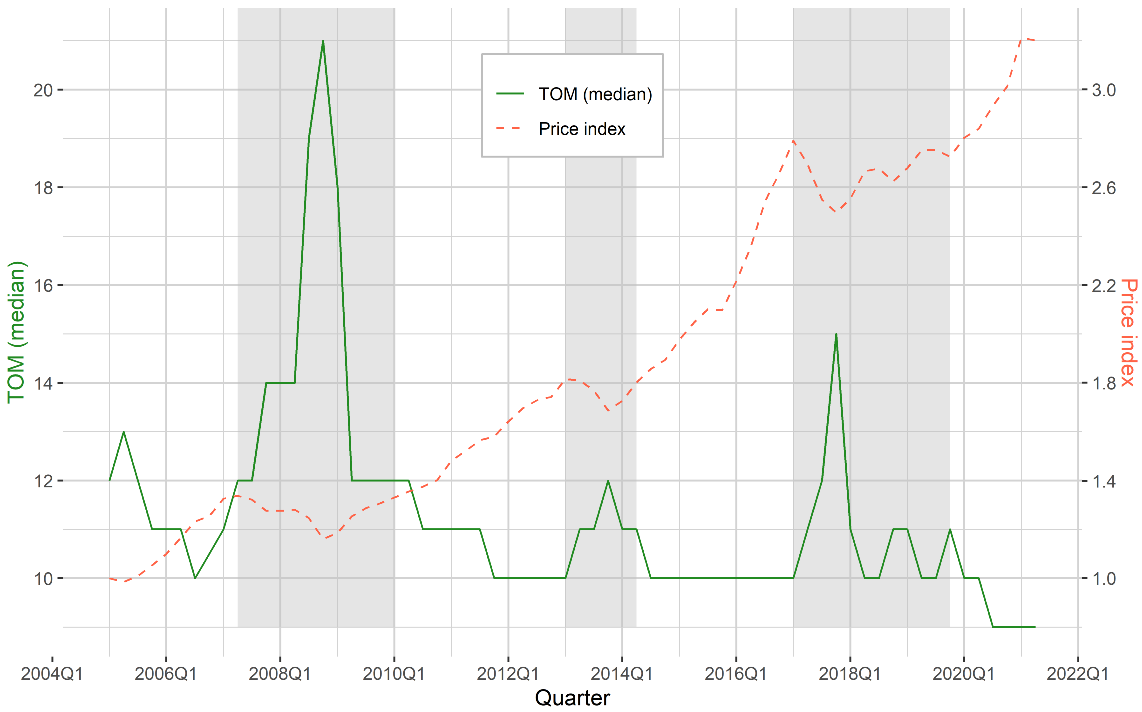

Notes: The figure presents a price index based on a hedonic prediction of house prices using year-by-quarter of sale dummies. Time-on-market is in days. The shaded areas begin at peaks and end when the index returns to the peak-level.

Table D.1: Data cleaning

Step	# of sellers/buyers	# of transactions
Initial	323,261	113,308
Drop co-ops before 2007	321,796	112,756
Remove duplicate rows	321,784	112,756
Remove transactions with missing list price	321,269	112,569
Remove transactions with less than 2 weeks between listings	321,076	112,496
Remove transactions with less than 2 weeks holding time	321,066	112,491
Remove holiday properties	321,020	112,475
To repeated: keep times sold>1	220,163	77,830
To repeated: matched buyers with sellers	194,741	68,472
Removing duplicate sellers and buyers	194,740	68,472
Reducing to one seller and one buyer	134,466	68,472
Reducing to one seller	66,603	66,603
Excluding first sale	41,019	41,019
Restrict to 2005-2020	37,401	37,401
Final trimmed Hedonic sample	32,044	32,044
Final trimmed Appraisal sample	16,111	16,111

Notes: The table reports the data cleaning process in detail, reporting numbers for Oslo. The column *# of sellers/buyers* is the number of rows in the sample, and *# of transactions* is the number of unique transactions.

Figure D.2: Annual observations

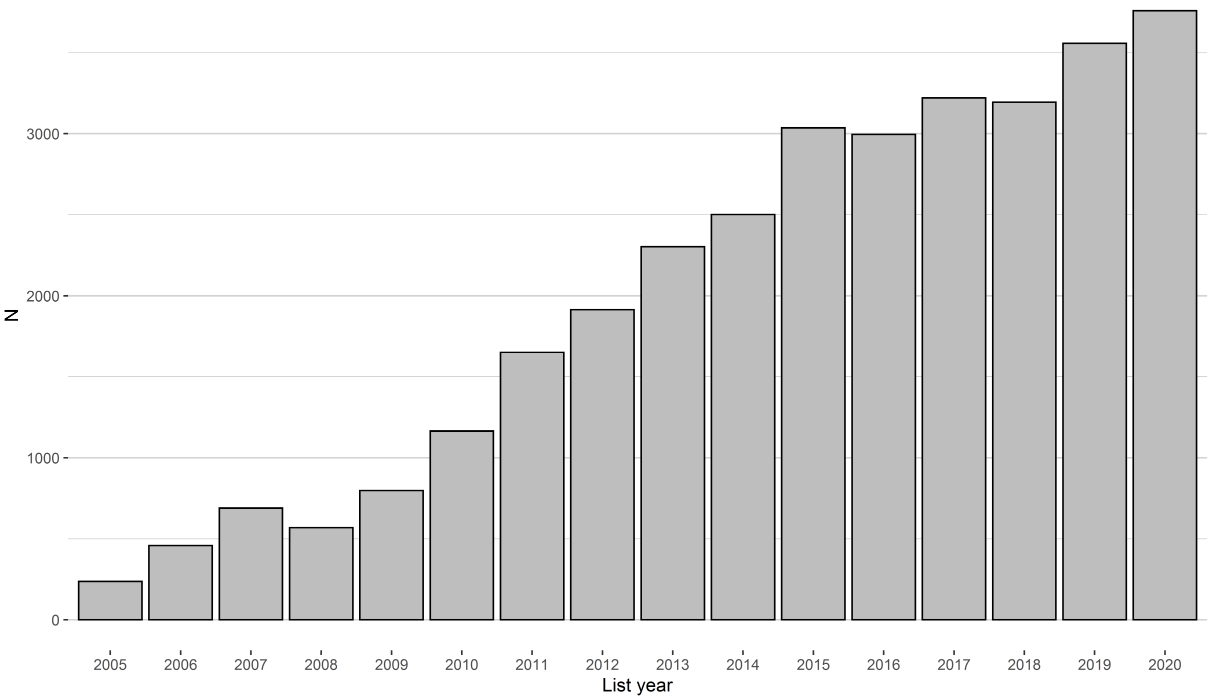

(A) Hedonic sample

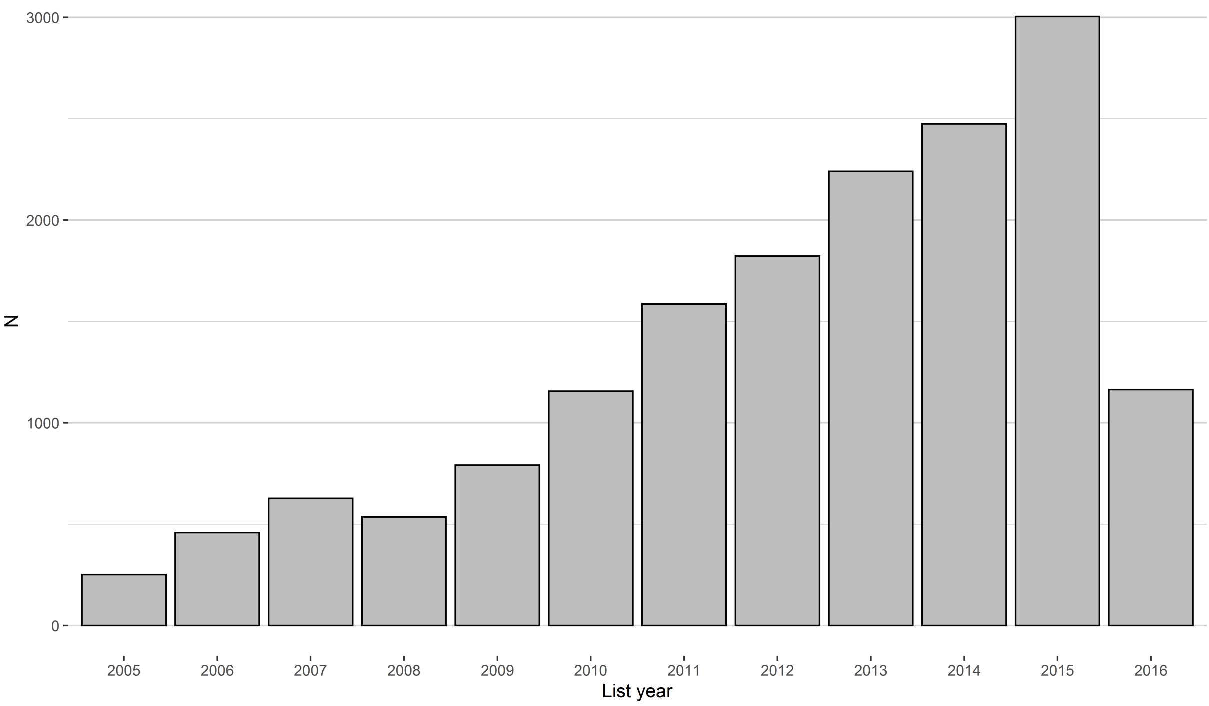

(B) Appraisal ample

Notes: The figure presents the frequencies of observations for each of the two samples of repeated purchases-to-listings. Panel A presents the frequencies for the hedonic sample, and panel B for the appraisal sample.

Table D.2: More summary statistics

Variable	1st Qu.	Median	Mean	3rd Qu.
(A) *Non-repeated sample, not trimmed (N=112,475, Dec 1999–Oct 2021)*				
List price (MNOK)	2.20	3.16	3.81	4.49
Selling price (MNOK)	2.35	3.30	3.96	4.65
Appraisal value (MNOK)	1.94	2.56	3.13	3.65
Size (m^2)	51	66	78	90
TOM (days)	9	11	24	19
Apartment (%)			85.68	
Non-co-op (%)			59.27	
(B) *Repeated sample, not trimmed (N=37,401, Jan 2005–Dec 2020)*				
List price (MNOK)	2.52	3.37	3.89	4.50
Selling price (MNOK)	2.69	3.50	4.04	4.67
Appraisal value (MNOK)	2.14	2.74	3.23	3.78
Size (m^2)	48	63	70	82
TOM (days)	9	10	22	17
Holding time (weeks)	125	197	226	294
Apartment (%)			90.68	
Non-co-op (%)			58.92	

Notes: The table reports summary statistics for the non-repeated sample after the initial cleaning, and the repeated sample before trimming. Prices are reported in millions of Norwegian kroner, and include common debt. Holding time is the number of weeks between previous date of sale and the date of listing. TOM denotes the time-on-market measured as the number of days between the listing and the sale.

Table D.3: Residual signs for zero and positive appraisal-based prospective losses

(A) $LOSS_{ist}^{AV} = 0$

		t: -	t: +
s	-	0.353	0.186
s	+	0.099	0.363

(B) $LOSS_{ist}^{AV} > 0$

		t: -	t: +
s	-	0.211	0.032
s	+	0.171	0.585

(C) Shares ($e_{i,\delta\in\{s,t\}} > 0$)

	s	t
$LOSS_{ist}^{AV} = 0$	0.461	0.549
$LOSS_{ist}^{AV} > 0$	0.755	0.617

Notes: The table reports the share of observations with positive or negative sign on the residual from the hedonic estimations. Panel A shows the shares conditional on zero prospective losses, which substitutes μ_{it} with P_{it}^{AV}. Panel B shows the shares conditional on positive prospective losses. Panel C shows the overall shares of positive residuals for the two periods.

Figure D.3: Price index and number of appraisal-based prospective losses

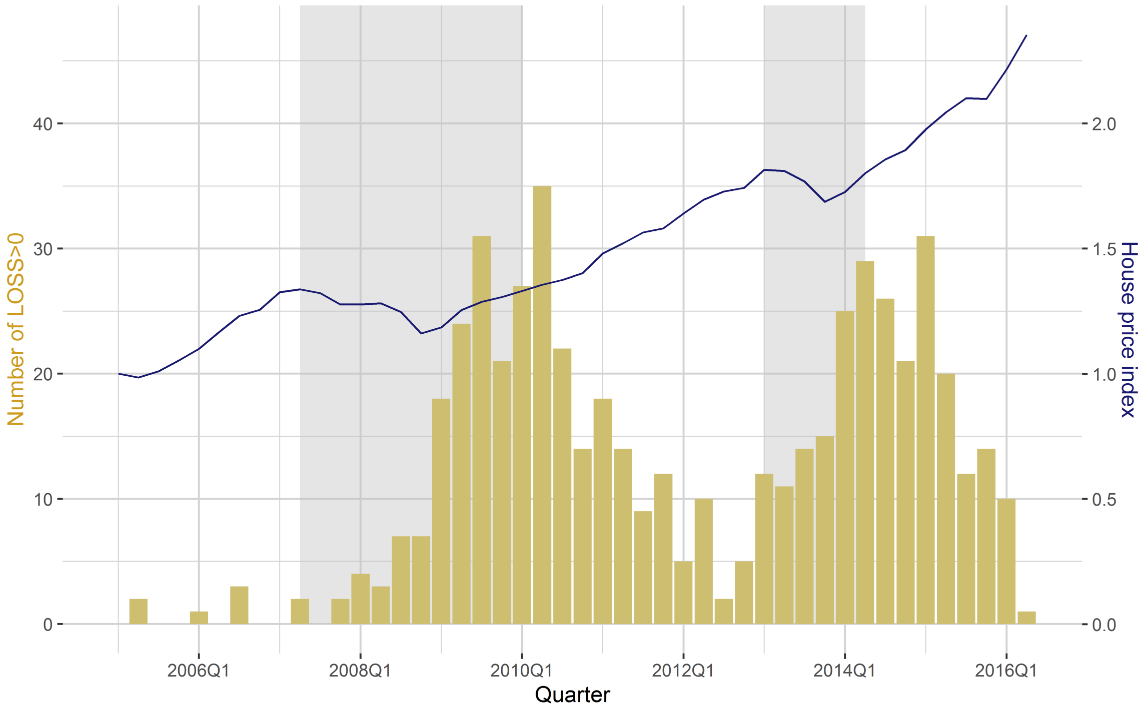

Notes: The figure presents a price index based on a hedonic prediction of house prices using year-by-quarter of sale dummies. The prospective losses use appraisal values as substitutes for price expectations. The shaded areas start at peaks and end when the index is back at the peak-level.

Table D.4: Pooling results for two-times repeated purchases-to-listings

	(1)	(2)	(3)	(4)
LOSS	0.918***	0.506***	−0.022	−0.009
	(0.072)	(0.087)	(0.066)	(0.065)
GAIN	0.341***	0.094**	0.031***	0.033***
	(0.021)	(0.041)	(0.005)	(0.005)
Substitute for μ_{it}	$\hat{P}_{it}$	$\hat{P}_{it}$	P_{it}^{AV}	P_{it}^{AV}
Residual control	No	Yes	No	Yes
N	10,324	10,324	4,026	4,026
Adj. R sq.	0.953	0.961	0.995	0.995
LOSS>0 (% of N)	7.594	7.594	4.223	4.223
VIF LOSS	1.135	1.194	1.134	1.076
VIF GAIN	1.768	3.452	1.514	1.732
F-stat.(LOSS=GAIN)	60.842	43.235	0.654	0.413
p-value(LOSS=GAIN)	<0.001	<0.001	0.419	0.520
LOSS-GAIN	0.577	0.412	-0.052	-0.041

Notes: The table presents results of pooling using balanced two-times repeated observations subsamples of purchases-to-listings. The dependent variable is log(List price). Predictions are of log(Selling price). The covariates include the log of holding time in weeks, DTV, the constant term for the pooling estimates, the predicted price, the log appraisal value, and the residual control (lagged selling price/appraisal value difference), which are omitted from the table. Standard errors are calculated using a wild bootstrap (R=1,000) and clustered on list year and 3-digit zip codes. Significance: * p<0.1, ** p<0.05, *** p<0.01.

Table D.5: Unit fixed effects - three-times repeated purchases-to-listings

	Pooling		Unit FE		Pooling		Unit FE	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
LOSS	0.719***	0.408***	0.306***	0.350***	−0.0003	0.010	−0.168*	−0.122
	(0.094)	(0.068)	(0.064)	(0.073)	(0.094)	(0.092)	(0.094)	(0.095)
GAIN	0.377***	0.156***	0.140***	0.173***	0.027***	0.029***	0.019***	0.026***
	(0.027)	(0.042)	(0.022)	(0.031)	(0.007)	(0.007)	(0.006)	(0.006)
Substitute for μ_{it}	$\hat{P}_{it}$	$\hat{P}_{it}$	$\hat{P}_{it}$	$\hat{P}_{it}$	P_{it}^{AV}	P_{it}^{AV}	P_{it}^{AV}	P_{it}^{AV}
Residual control	No	Yes	No	Yes	No	Yes	No	Yes
N	4,371	4,371	4,371	4,371	897	897	897	897
Adj. R sq.	0.960	0.965	0.950	0.950	0.996	0.996	0.981	0.981
LOSS>0 (% of N)	8.945	8.945	8.419	8.419	5.463	5.463	5.128	5.128
VIF LOSS	1.153	1.206	1.126	1.150	1.093	1.108	1.075	1.089
VIF GAIN	1.621	3.122	1.851	2.612	1.393	1.586	1.352	1.432
F-stat.(LOSS=GAIN)	11.836	16.235	5.574	6.183	0.080	0.043	3.744	2.260
p-value(LOSS=GAIN)	0.001	<0.001	0.018	0.013	0.778	0.836	0.053	0.133
LOSS-GAIN	0.342	0.252	0.166	0.177	-0.027	-0.020	-0.187	-0.147

Notes: The table presents results of pooling and unit fixed effect using balanced three-times repeated observations subsamples of purchases-to-listings. Columns (1), (2), (5) and (6) are pooling estimates, and (3), (4), (7) and (8) are unit fixed effects (or first difference). Moreover, (1)-(4) are hedonic-based estimates while (5)-(8) are appraisal-based. The dependent variable is log(List price). Predictions are of log(Selling price). The covariates include the log of holding time in weeks, DTV, the constant term for the pooling estimates, the predicted price, the log appraisal value, and the residual control (lagged selling price/appraisal value difference), which are omitted from the table. Standard errors are calculated using a wild bootstrap (R=1,000) and clustered on list year and 3-digit zip codes. Significance: * p<0.1, ** p<0.05, *** p<0.01.

Table D.6: Aggregate prospective terms

	(1)	(2)	(3)	(4)
LOSS	0.329^{***}	0.507^{***}	-0.091^{*}	-0.098^{*}
	(0.085)	(0.085)	(0.048)	(0.053)
GAIN	0.020	0.024	-0.022^{***}	-0.023^{***}
	(0.018)	(0.018)	(0.005)	(0.005)
Outside price	$\hat{P}_{it}$	$\hat{P}_{it}$	P_{it}^{AV}	P_{it}^{AV}
Residual control	No	Yes	No	Yes
N	31,536	30,450	15,647	13,981
Adj. R sq.	0.935	0.960	0.996	0.996
LOSS>0 (% of N)	2.549	2.631	2.870	2.847
VIF LOSS	1.045	1.047	1.052	1.052
VIF GAIN	2.590	2.594	2.293	2.287
F-stat.(LOSS=GAIN)	14.690	35.979	2.066	2.007
p-value(LOSS=GAIN)	<0.001	<0.001	0.151	0.157
LOSS-GAIN	0.309	0.483	-0.069	-0.075

Notes: The table presents results of aggregated prospective term, meaning $LOSS_{ist} = (\delta_s - \delta_t)^+$ and $GAIN_{ist} = (\delta_s - \delta_t)^-$. Columns (1) and (2) are hedonic-based estimates, and (3) and (4) are appraisal-based estimates. The dependent variable is log(List price). Predictions are of log(Selling price). The covariates include the log of holding time in weeks, DTV, the constant term, the predicted price, the log appraisal value, and the residual control (lagged selling price/appraisal value difference), which are omitted from the table. Standard errors are calculated using a wild bootstrap (R=1,000) and clustered on list year and 3-digit zip codes. Significance: * p<0.1, ** p<0.05, *** p<0.01.

Figure D.4: Effects of adding random noise, with lagged residual

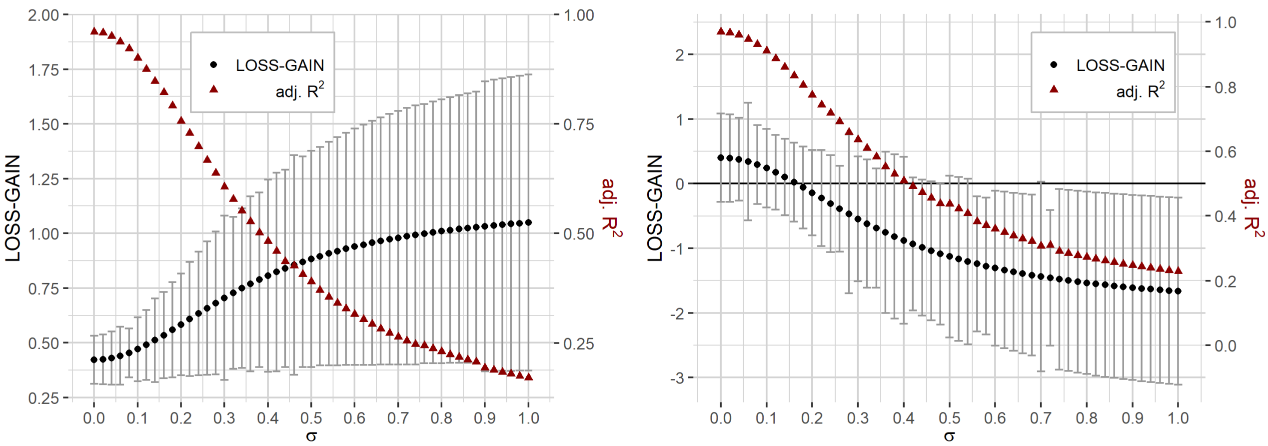

(A) Hedonic-based LOSS-GAIN

(B) Appraisal-based LOSS-GAIN

Notes: The figure presents the difference between prospective loss and gain, a 95 percent confidence interval, and the adjusted R^2 when adding noise $\xi \sim N(0, \sigma^2)$ with different standard deviations σ to the predicted price $\hat{P}_{it}$. For each $\sigma > 0$, the procedure of drawing ξ and estimating the coefficient difference is repeated 1,001 times and the median coefficient difference between loss and gain, with the corresponding confidence interval, is chosen to be plotted. This is done because the loss-gain distribution is normal. Panel A shows the hedonic-based loss and gain difference, meaning using the hedonic price prediction as the inside price substitute. Panel B show the appraisal-based loss and gain difference. Standard errors are clustered on list year and 3-digit zip codes.

Table D.7: Splitting the sample into short and long holding time

	Short hold		Long hold	
	(1)	(2)	(3)	(4)
LOSS	0.727***	0.417***	0.809***	0.365***
	(0.060)	(0.066)	(0.158)	(0.119)
GAIN	0.400***	0.171***	0.367***	0.045
	(0.029)	(0.055)	(0.027)	(0.063)
Residual control	No	Yes	No	Yes
N	16,062	15,516	16,034	14,979
Adj. R sq.	0.959	0.965	0.950	0.956
LOSS>0 (% of N)	12.794	13.019	0.823	0.855
VIF LOSS	1.195	1.269	1.026	1.034
VIF GAIN	1.273	2.524	1.777	5.394
F-stat.(LOSS=GAIN)	24.837	32.857	6.651	7.969
p-value(LOSS=GAIN)	<0.001	<0.001	0.010	0.005
LOSS-GAIN	0.327	0.246	0.442	0.320

Notes: The table presents results of separating the data into a short and a long holding time subsample, separated based on the median holding time (198 weeks). Columns (1) and (2) are results from the short holding time subsample, and (3) and (4) are from the long holding time subsample. The dependent variable is log(List price). Predictions are of log(Selling price). The covariates include the log of holding time in weeks, DTV, the constant term, the predicted price, and the residual control, which are omitted from the table. Standard errors are wild bootstrapped (R=1,000) and clustered on list year and 3-digit zip codes. Significance: * p<0.1, ** p<0.05, *** p<0.01.

Table D.8: Log appraisal value as dependent variable

	(1)	(2)
LOSS	0.964***	0.632***
	(0.048)	(0.050)
GAIN	0.415***	0.202***
	(0.019)	(0.036)
Residual control	No	Yes
N	15,901	14,583
Adj. R sq.	0.949	0.955
LOSS>0 (% of N)	8.471	8.969
VIF LOSS	1.122	1.266
VIF GAIN	1.560	3.936
F-stat.(LOSS=GAIN)	124.807	125.893
p-value(LOSS=GAIN)	<0.001	<0.001
LOSS-GAIN	0.550	0.429

Notes: The table presents results from using log(Appraisal value) as dependent variable, predictions have log(Selling price) as dependent variable. The difference in number of observations of the lower bound estimates presented here compared to Table 4 comes from the balancing not done here. In the baseline lower bound estimation, I require both that the lagged appraisal value and the lagged hedonic residual is available, while the estimations presented here only require the lagged hedonic residuals. The covariates include the log of holding time in weeks, DTV, the constant term, the predicted price, and the residual control, which are omitted from the table. Standard errors are wild bootstrapped (R=1,000) and clustered on list year and 3-digit zip codes. Significance: * $p<0.1$, ** $p<0.05$, *** $p<0.01$.

Table D.9: Removing the top most influencing observations based on Cook's distance

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
LOSS	0.907***	0.409***	0.046***	0.036**	0.913***	0.408***	0.040*	0.024
	(0.045)	(0.057)	(0.017)	(0.018)	(0.048)	(0.056)	(0.021)	(0.024)
GAIN	0.312***	0.039	0.027***	0.027***	0.313***	0.038	0.026***	0.027***
	(0.017)	(0.032)	(0.003)	(0.004)	(0.017)	(0.032)	(0.003)	(0.004)
Substitute for μ_{it}	$\hat{P}_{it}$	$\hat{P}_{it}$	P_{it}^{AV}	P_{it}^{AV}	$\hat{P}_{it}$	$\hat{P}_{it}$	P_{it}^{AV}	P_{it}^{AV}
Residual control	No	Yes	No	Yes	No	Yes	No	Yes
Removed influencers	Top 5	Top 5	Top 5	Top 5	Top 10	Top 10	Top 10	Top 10
N	34,813	31,803	18,161	16,190	34,808	31,798	18,156	16,185
Adj. R sq.	0.950	0.960	0.995	0.995	0.951	0.960	0.995	0.995
LOSS>0 (% of N)	8.382	8.792	4.240	4.373	8.377	8.787	4.219	4.356
VIF LOSS	1.110	1.249	1.041	1.047	1.110	1.249	1.041	1.056
VIF GAIN	1.780	3.634	1.581	1.628	1.783	3.634	1.581	1.711
F-stat.(LOSS=GAIN)	202.781	118.759	1.529	0.290	174.351	123.607	0.451	0.013
p-value(LOSS=GAIN)	<0.001	<0.001	0.216	0.590	<0.001	<0.001	0.502	0.910
LOSS-GAIN	0.595	0.370	0.019	0.009	0.600	0.369	0.014	-0.003

Notes: The table presents results from non-trimmed samples but removing the top influencing observations. Columns (1) to (4) removes the top five most influencing observations, while columns (5) to (8) removes the top ten. The dependent variable is log(List price). Predictions are of log(Selling price). The covariates include the log of holding time in weeks, DTV, the constant term, the predicted price, and the residual control, which are omitted from the table. Standard errors are calculated using a wild bootstrap (R=1,000) and clustered on list year and 3-digit zip codes. Significance: * p<0.1, ** p<0.05, *** p<0.01.

Figure D.5: Cook's distance

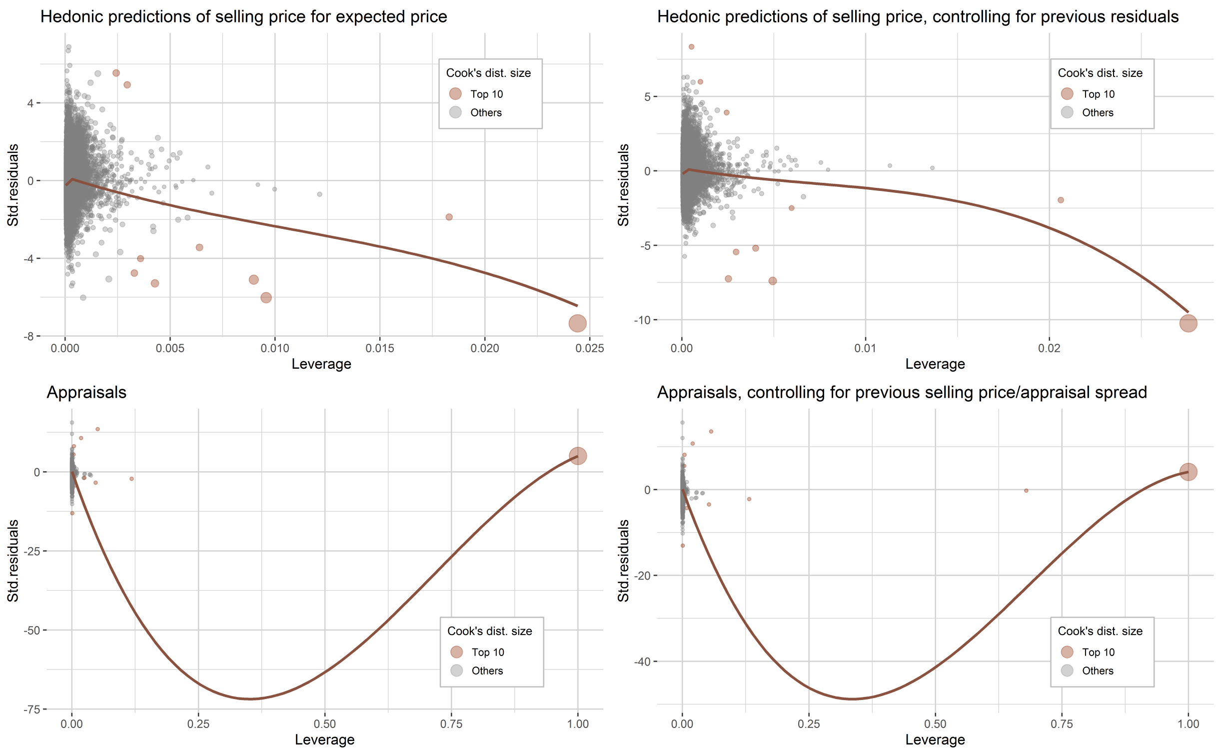

Notes: The figure presents Cook's distance plots for the models estimated in table D.9. The two upper plots are for the hedonic-based models, while the two lower plots are for the appraisal-based models. The y-axis are standardized residuals, and the x-axis is the diagonal of the projection matrix.

Figure D.6: Hockey stick

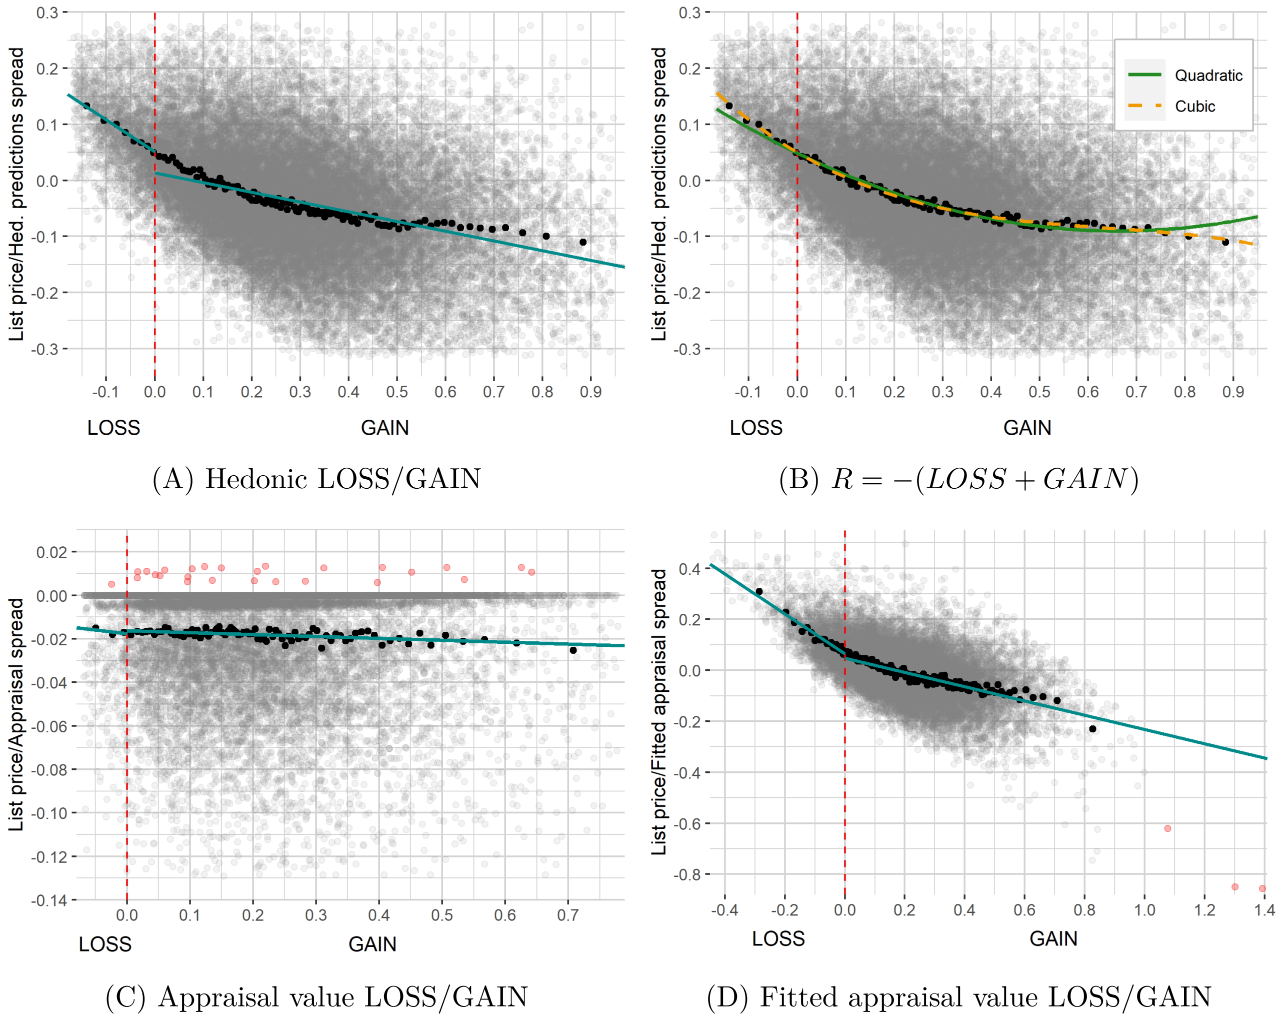

(A) Hedonic LOSS/GAIN

(B) $R = -(LOSS + GAIN)$

(C) Appraisal value LOSS/GAIN

(D) Fitted appraisal value LOSS/GAIN

Notes: The figure presents the unconditional hockey stick pattern between the list price premium and the prospective variables. The x-axis is the difference between the price expectation and the previous selling price, $R = -(LOSS + GAIN)$, thus having opposite sign of those used in the regressions. Solid black points are the means of the spread within bins of R. Panel A and B show the hedonic-based relationship, with panel B showing that using appropriate functional forms make the inclusion of prospective loss as a separate variable irrelevant for the fit. Panel C shows the appraisal-based relationship, where observations with the list price above the appraisal value are highlighted. Panel D shows the fitted appraisal-based relationship, where observations with prospective gains above 1 are highlighted. Highlighted observations are colored red.

Table D.10: Hedonic-based model: sensitivity to functional forms

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
LOSS	0.649***	0.296***	0.588***	0.255***	0.095	0.042	−0.095	−0.091
	(0.057)	(0.047)	(0.052)	(0.046)	(0.082)	(0.047)	(0.102)	(0.069)
GAIN	0.521***	0.229***	0.605***	0.289***				
	(0.046)	(0.063)	(0.062)	(0.070)				
$GAIN^2$	0.275***	0.213***	0.548***	0.409***				
	(0.051)	(0.042)	(0.134)	(0.090)				
$GAIN^3$			0.227***	0.166***				
			(0.082)	(0.056)				
RF					−0.522***	−0.229***	−0.611***	−0.293***
					(0.046)	(0.063)	(0.064)	(0.070)
RF^2					0.276***	0.214***	0.563***	0.419***
					(0.052)	(0.042)	(0.137)	(0.091)
RF^3							−0.238***	−0.173***
							(0.083)	(0.056)
Substitute for μ_{it}	$\hat{P}_{it}$	$\hat{P}_{it}$	$\hat{P}_{it}$	$\hat{P}_{it}$	$\hat{P}_{it}$	$\hat{P}_{it}$	$\hat{P}_{it}$	$\hat{P}_{it}$
Residual control	No	Yes	No	Yes	No	Yes	No	Yes
N	32,044	30,450	32,044	30,450	32,044	30,450	32,044	30,450
Adj. R sq.	0.954	0.961	0.954	0.961	0.954	0.961	0.954	0.961
LOSS>0 (% of N)	6.856	7.087	6.856	7.087	6.856	7.087	6.856	7.087
F-stat.(LOSS=GAIN)	2.681	2.402	0.037	0.335				
p-value(LOSS=GAIN)	0.102	0.121	0.847	0.563				
LOSS-GAIN	0.128	0.068	-0.017	-0.034				

Notes: The table presents results from including different functional forms of prospective gain and a reference dependence variable $R = LOSS + GAIN$. The dependent variable is log(List price). Predictions are of log(Selling price). The covariates include the log of holding time in weeks, DTV, the constant term, the predicted price, and the residual control, which are omitted from the table. Standard errors are calculated using a wild bootstrap (R=1,000) and clustered on list year and 3-digit zip codes. Significance: * p<0.1, ** p<0.05, *** p<0.01.

3 Here Comes the Sun: The Effect of Sunshine on Home Prices

Here Comes the Sun: The Effect of Sunshine on Home Prices*

Andreas Eidspjeld Eriksen† Cloé Garnache‡

November 11, 2024

Abstract

Using high-frequency data for the Norwegian housing market, we show that sunshine affects residential real estate prices. In particular, we find evidence that sunshine affects asset price formation primarily during the days of public showings of for-sales homes, rather than during the final bidding process. Furthermore, we find no evidence that sunshine affects buyer interest, including the number of participants at showings and the number of people making offers. These results suggest that the primary channel through which sunshine impacts households' beliefs and asset prices is during the initial information processing stage.

Keywords: *Asset pricing; Price formation; Housing market; Sunshine; Cloud cover; Mood*
JEL classification: *D84; G12; G14; G41; R21; R31*

1 Introduction

Empirical evidence shows that local cloud cover is negatively correlated with stock market returns (e.g., Saunders, 1993; Hirshleifer and Shumway, 2003; Goetzmann and Zhu, 2005). Institutional investors, commonly perceived as sophisticated agents, are not exempt from this effect – they are more likely to perceive stocks as overpriced and sell during cloudy days (Goetzmann et al., 2015). The effect of sunshine on mood and behavior has

*We are grateful to Bjørnar Karlsen Kivedal for helpful comments. We also want to thank Eiendomsverdi AS and DNB Eiendom for providing data.

†School of Economics and Business, Norwegian University of Life Sciences and Housing Lab, Oslo Metropolitan University. E-mail: andrease@oslomet.no

‡Housing Lab, Oslo Metropolitan University and Dept. of Economics, University of Oslo. E-mail: cloe.garnache@econ.uio.no

long been documented in psychology studies.[1] An open question is the channel through which sunshine-induced mood impacts investors' belief and price formation. Does sunshine influence mostly information processing during the initial assessment or during the decision-making stage?[2]

This paper first examines whether sunshine, or more specifically its counterpart, cloud cover, influences the residential real estate market. On the one hand, because purchasing a home represents a large financial decision, one may expect households to behave rationally and not be influenced by weather-induced mood. On the other hand, findings from psychology suggest households evaluating the expected stream of future cash flows from a home may be particularly susceptible to weather-induced mood biases.[3] Since housing is often the largest and most significant asset for households (Campbell and Cocco, 2007), any biases in how households form price expectations could have substantial impacts on their savings and debt. Moreover, the misattribution of mood as information may distort aggregate residential real estate prices, carrying broader macroeconomic implications (Leamer, 2007).

Second, to identify the channels through which cloud cover impact households' price formation, this paper leverages cloud cover variation during different stages of the home sales process, including during showings and bids submission. Identifying the impact of cloud cover can be challenging, as it is often unclear when prospective buyers visit a home and submit an offer. To address this issue, we use data from the Norwegian housing market, where the sales process is highly standardized: homes usually have a very short time-on-market (typically less than two weeks), including one or two public showings, followed by an online auction a day or two later. A home is recorded as sold once the seller accepts a bid. This structured process allows us to pinpoint the timing of key stages of the sales process, in particular, the showings and the bid submission, thus, enabling us to measure cloud cover on these relevant days. Because other weather conditions may affect prospective buyers' perceptions and/or interact with cloud cover-induced mood, we

[1] The effect of sunshine on beliefs and behavior has been documented in various settings, such as daily mood (Denissen et al., 2008), life satisfaction (Schwarz and Clore, 1983), and altruistic behavior (Cunningham, 1979).

[2] Psychology studies document that affective states may affect memory encoding, recall, and decision making (e.g., Clore et al., 1994; Schwarz and Clore, 2007; Schwarz, 2011; Forgas, 2017, for reviews). For example, Schwarz and Clore (1983) find that individuals may mistakenly attribute mood induced by sunshine as informative cues when forming beliefs, Forgas (1995) shows that affect is important for cognition when faced with a new and demanding task, while Cunningham (1979) and Loewenstein (2000) suggest that emotions may affect decision-making at the moment of acting.

[3] Mood is especially relevant as households typically visit a home during a showing and make an offer days later relying on their potentially mood-tainted memory of the home. Furthermore, the task of assessing a home's stream of future cash flows is complex and may be prone to multiple mood-related biases. For instance, affective states have been shown to affect how individuals predict future states of the world (Johnson and Tversky, 1983), and make assessments in situations with limited information (Schwarz and Clore, 2007, p. 7).

control for an array of weather variables and interaction effects.[4]

Our findings are two-fold. First, findings suggest that home prices are negatively affected by cloud cover. Using our preferred model specification with housing unit fixed effects, a standard deviation increase in cloud cover from the mean at the showings, reduces home prices by 0.83% (p-value$<$0.05), at the sample means, corresponding to approximately $2,908 for the median home in our sample. This effect is equivalent to a 6.57% increase in the mortgage rate during our time period (2015-2021), based on the mortgage rate coefficient estimate from the corresponding model. Results are robust to a series of alternative weather specifications, inclusion of alternative fixed effects, and sample choices. A placebo test with random draws of weather conditions indicates no economically meaningful effect of cloud cover.

Second, we find no evidence that cloud cover during the day of the sale influences real estate prices. We also find no evidence that cloud cover affects market liquidity, including the number of prospective buyers attending showings, the number of bidders, and the sell-ask spread. Our results suggest that cloud cover primarily influences information processing during the initial assessment stage (public showings of for-sales homes), rather than during the final decision-making stage (the auction) when prospective buyers might adjust their initial valuation.

This paper contributes to three strands of literature. First, it relates to the finance literature that explores the channels through which weather-induced mood influences economic agents' beliefs and price formation. Goetzmann et al. (2015) shows that cloud cover influences institutional investors' beliefs about the stock market, with investors being more optimist about the pricing of individual stocks and the Dow Jones Industrial Average when the sun shines. Our study investigates how cloud cover-induced mood affects the different stages of the price formation process, from initial information processing to final assessment and decision-making.[5]

Second, our study adds to the literature on the impact of cloud cover on investor behavior and stock market index returns. For example, Saunders (1993) estimates a negative relationship between cloud cover on Wall Street – plausibly unrelated to firm

[4]Because the winter season in Norway accounts for a small share of home sales (17%), is characterized by few hours of daylight (7.12 hours on average in the sample) and, thus, may not be representative of other residential real estate markets, we exclude winter from the main analysis (see Appendix C).

[5]Beyond the effects of cloud cover, Kamstra et al. (2003) demonstrate that shorter daylight hours in winter affect stock market returns. Additionally, Kamstra et al. (2014) present a theoretical model incorporating seasonally varying risk preferences and elasticity of intertemporal substitution to explain fluctuations in Treasury and equity returns.

Bodoh-Creed (2020) develops a model exploring the effect of mood-congruent memory on asset prices, where the agent's mood serves as an informational cue stored in long-term memory. The model predicts that more sophisticated market participants may be more vulnerable to mood-related biases. Our study does not confirm this prediction as we find that more experienced and wealthier buyers are less influenced by cloud cover than first-time buyers and less affluent households.

fundamentals – and stock returns at the New York City stock exchanges. Similarly, Hirshleifer and Shumway (2003) find a similar relationship between local cloud cover and stock market index returns across a broad range of countries. Their findings suggest that, globally, cloud cover decreases investor optimism and their propensity to buy stocks. Goetzmann and Zhu (2005) provide evidence indicating that cloud cover in New York City influences New York Stock Exchange market. Yet, they do not find evidence that cloud cover affects retail investors' propensity to sell or buy. Goetzmann et al. (2015) show that cloud cover impacts institutional investors' trading behavior, though their observations are limited to quarterly investor positions. Our study extends these findings to a new asset class – residential real estate – which, in addition to being typically households' largest financial asset, allows us to precisely isolate the different stages of information processing during the decision-making process.

Third, to our knowledge, this is the first study investigating the effect of cloud cover-induced mood in the residential real estate market. While one might expect households to act rationally and avoid mood-driven biases when making large stake decisions such as purchasing a home, our findings suggest that cloud cover distorts residential real estate prices. The closest related study is Gourley (2021), who investigates the impact of temperature and precipitation, without controlling for cloud cover, on home prices in Jefferson County, Colorado, using monthly weather averages during the month preceding the closing date.[6] The purchase agreement may take place a few days to up to several months prior to the closing date because several steps must occur after the purchase agreement, including the home inspection and mortgage application process. It is thus difficult to pinpoint when the purchase agreement or even the showings take place. In contrast, our setting allows us to precisely isolate the key stages of the price formation process.

2 Data

2.1 Data and study sample

This section briefly describes data sources, variables construction, and motivates the choice of the study sample. Additional details are provided in Appendix A.

Our study uses data from several sources. First, we obtain listings for for-sale homes and sales transactions from Eiendomsverdi AS. The data consist of sell price, attributes, listing and sale dates, and zip code, and cover the largest municipalities in the south-eastern part of Norway in the period 2015-2021. Sell prices are deflated to 2015 prices and converted to United States dollars (USD) using the exchange rate prevailing on January

[6] Variables with strong seasonal dependence, such as monthly averages, are prone to proxy for seasonality, as discussed by, e.g., Jacobsen and Marquering (2008) and Kelly and Meschke (2010).

2, 2023 (NOK/USD=9.8413). All transactions are between private persons and consist of second-hand on-market sales. There are two main ownership types in Norway, co-ops and non-co-ops, and there are four different housing unit types: apartments, detached, semi-detached, and row houses.[7] Other attributes include the size of the living area and the lot size, both in square meters.

Second, we retrieve three weather-related data sets from the Copernicus Climate Change Service (Copernicus Climate Change Service, 2024a,c,b): (i) daily daytime cloud cover with a median grid cell size of 336 square kilometers, measured on a scale from 0 to 1, where 0 is clear sky and 1 is fully cloud covered sky (Karlsson et al., 2023), (ii) daily accumulated precipitation and maximum temperature at a one-by-one kilometer resolution (Tveito et al., 2000, 2005), and (iii) daily forecasts of wind speed based on the ERA5 model with a median grid cell size of 100 square kilometers (Muñoz Sabater, 2019). Because daylight hours likely interact with cloud cover perceptions, we define a variable 'hours of night' (hereafter HoN), which we calculate using a similar method as the one proposed in Kamstra et al. (2003). All weather and HoN variables are aggregated to the zip code level. Following Goetzmann et al. (2015), we construct (monthly) seasonal, local controls by averaging each of the weather and HoN variable across all years (2015-2023) for each zip code and month between March and November. All weather and HoN variables, including seasonal variables, are transformed using natural logarithm.

Third, we obtain buyer characteristics and monthly aggregated mortgage rates reported by consumer banks from Statistics Norway. Buyer characteristics include (end-of-year) income, debt, wealth, household type, number of children in the household, number of adults in the household, education level, age, and individuals' gender in the period 2014-2019.

Fourth, we use a supplementary data set containing the number of showing participants from public showings and number of bidders for sales transactions between August 2018 and June 2023 from the realtor firm, DNB Eiendom. DNB Eiendom is the second largest realtor firm in Norway with about 20 percent market share.

Our sample contains transactions from the 13 largest municipalities (based on 2021 population), in the south-east of Norway (Figure 1). Climate in Norway differs across regions. For instance, the west coast of Norway is known to be heavily exposed to wind and precipitation, while the North, located above the arctic circle, is cold and with extreme daylight variation throughout the year. For this study, we focus on municipalities in the south-east of Norway as they experience milder and relatively homogeneous weather conditions. Our study excludes the winter season, defined as December through February,

[7] The co-ops differ from non-co-ops mostly in the ownership structure, in which owners of housing units in co-ops owns their home indirectly through a co-op association. Both of these ownership types are common in large cities, while in rural areas non-co-ops are predominant.

as the housing market slows down with 17% of the sales in the sample taking place during this period, and because daylight hours during this period are short (7.12 hours on average) in the study area (located between latitudes 58 and 60), thus, limiting the representativity of this season for other real estate markets. Moreover, other factors may correlate with cloud cover and home prices in the winter, such as the snow cover and depth. As an extension, we estimate the main model for all four seasons in Appendix C. The overall effects become smaller, as expected from adding more noise. When allowing all weather variables to have a different slope in the winter season, cloud cover at showings has no effect on home prices in the winter, while the effects in the remaining seasons are similar to those in our main results.

Figure 1: Sample municipalities in the south-east of Norway

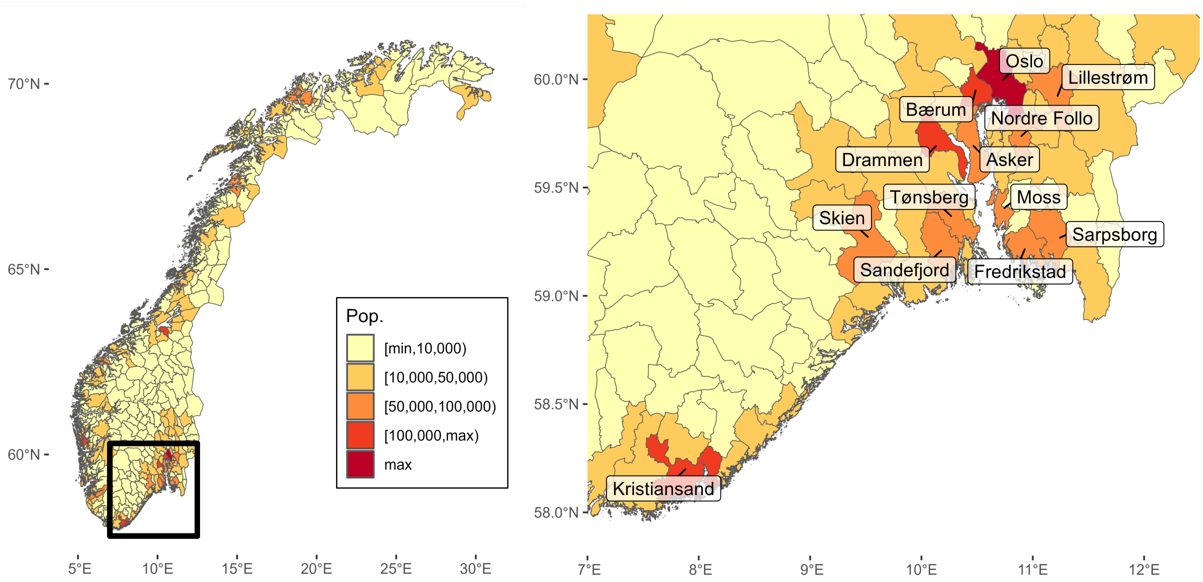

Notes: The figure shows the location of the municipalities in our sample and the 2021 population by municipality for Norway (left panel) and for the south-eastern part of Norway with the names of the municipalities included in our sample (right panel).

2.2 Cloud cover during the different stages of the sales process

The sales process in the Norwegian housing market is highly standardized, which enables us to recover the day of the sale and infer the likely dates of the showings relatively precisely. The sales process proceeds as follows: first, the housing unit is listed online. The for-sales listing features the dates for the public showings and for the auction. Second, prospective buyers visit the unit at the public showing(s). Third, the day after the last showing (as announced in the for-sales listing), prospective buyers submit bids and the unit is typically sold the same day.

Based on our supplementary data set with exact showing dates, 79% of sales happen one day after the last showing, while 85% of sales happen within 4 days after the first showing – even when conditioning on municipalities (Figure A.1). Therefore, to investigate the effect of cloud cover at the showings, we measure cloud cover and other weather variables as the moving average of the four days prior to the date of sale (MA). That is, the moving average lagged one day from the sales date.

To examine the effect of cloud cover during the final stage of the sales process, i.e., the bid submission during the auction, we measure cloud cover and other weather variables during the day of the sale. Summary statistics including weather at the day of the sale (DOS) and the (one day lagged) moving averages are presented in Table 1.

3 Empirical framework

This section introduces two main specifications for estimating the effect of cloud cover-induced mood on housing market outcomes. First, we propose a specification to estimate the effect of cloud cover at the showings on sell prices. Second, we propose a two-step approach to disentangle the effect of cloud cover at the showings and at the day of the sale.

3.1 Effect of cloud cover at public showings

Our main regression specification estimates the effect of cloud cover at the showings on sell prices. The regression specification is

$$\begin{aligned} P_{it} = \alpha &+ \beta_1 \mathrm{CC}_{jt}^{MA} + \beta_2 \mathrm{CC}_{jt}^{MA} \times \mathrm{Prec}_{jt}^{MA} + \beta_3 \mathrm{CC}_{jt}^{MA} \times \mathrm{HoN}_{jt} \\ &+ \beta_4 \mathrm{Prec}_{jt}^{MA} + \beta_5 \mathrm{HoN}_{jt} + \beta_6 \mathrm{Prec}_{jt}^{MA} \times \mathrm{HoN}_{jt} \\ &+ \beta_7 \mathrm{PosMaxTemp}_{jt}^{MA} + \beta_8 \mathrm{NegMaxTemp}_{jt}^{MA} + \beta_9 \mathrm{Wind}_{jt}^{MA} \\ &+ \mathbf{W}_{jm} \beta_{\mathbf{W}} + \beta_{10} \mathrm{MortgageRate}_t + \mathbf{X}_i \beta_{\mathbf{X}} + \mu_{it} + Zip_j + \varepsilon_{it}, \end{aligned} \tag{1}$$

where P_{it} is the log of real sell price for unit i at the sale date t. The weather variables $w \in W$ represent the logs of the moving averages of the four days prior to the sale date t, denoted with superscript MA, in zip code j: CC_{jt}^{MA} is the average daytime cloud cover, Prec_{jt}^{MA} is the precipitation, $\mathrm{PosMaxTemp}_{jt}^{MA}$ is the maximum temperature if the moving average is positive and else zero, $\mathrm{NegMaxTemp}_{jt}^{MA}$ is the absolute value of the maximum temperature if the moving average is negative and else zero, and Wind_{jt}^{MA} is

Table 1: Summary statistics

N = 55,823	Mean	SD	25th pct.	Median	75th pct.
Home sales					
Sell price (in 1,000 USD)	411.20	211.30	274.63	350.16	483.87
Size (sqm.)	88.18	47.20	55.00	73.00	109.00
Age (years)	52.12	36.84	27.00	49.00	67.00
TOM (days)	10.57	3.30	8.00	10.00	11.00
Lot size (sqm.)	2,001.20	9,853.77	437.00	744.00	1,068.00
Apartments (share)	0.70				
Owner-occupiers (share)	0.60				
Mortgage rate	2.40	0.41	2.22	2.45	2.62
Across seasons					
Cloud cover (MA)	0.59	0.25	0.42	0.61	0.78
Precipitation (mm/day, MA)	2.85	3.87	0.06	1.28	4.15
Cloud cover (DOS)	0.59	0.38	0.19	0.69	0.98
Precipitation (mm/day, DOS)	3.03	6.99	0.00	0.02	2.44
Hours of night	10.05	3.52	6.93	9.67	12.63
Spring					
Cloud cover (MA)	0.63	0.25	0.48	0.66	0.82
Precipitation (mm/day, MA)	2.14	2.84	0.02	0.82	3.39
Cloud cover (DOS)	0.59	0.39	0.18	0.71	1.00
Precipitation (mm/day, DOS)	1.96	4.75	0.00	0.00	1.29
Hours of night	9.39	2.40	7.27	9.02	11.62
Summer					
Cloud cover (MA)	0.51	0.23	0.36	0.52	0.69
Precipitation (mm/day, MA)	3.05	3.75	0.11	1.64	4.59
Cloud cover (DOS)	0.50	0.36	0.13	0.51	0.86
Precipitation (mm/day, DOS)	3.02	7.09	0.00	0.01	1.90
Hours of night	6.99	1.73	5.47	5.88	8.79
Fall					
Cloud cover (MA)	0.63	0.26	0.44	0.66	0.86
Precipitation (mm/day, MA)	3.54	4.83	0.11	1.57	5.09
Cloud cover (DOS)	0.67	0.37	0.36	0.84	1.00
Precipitation (mm/day, DOS)	4.35	8.76	0.00	0.23	4.73
Hours of night	13.92	2.28	11.94	13.96	15.85

Notes: The table presents summary statistics of the main variables. Sell prices are deflated by the consumer price index (CPI) to 2015 prices, then converted to USD using the January 2, 2023 exchange rate (NOK/USD=9.8413). Size refers to the living area. TOM is the time-on-market, being the number of days between the listing and the sale. *MA* indicates the moving averages of the four days prior to the sale, hence, a one day lagged moving average. *DOS* indicates the day-of-the-sale. Lot size is summarized for non-apartments.

the instantaneous wind speed forecast in the early afternoon. HoN_{jt} is the log of the hours of night in the zip code j where unit i is located on sale date t.

Because the effect of cloud cover on prospective buyer behavior is likely to vary with the levels of precipitation and hours of night, we include interaction terms ($CC \times Prec$ and $CC \times HoN$). We further control for the interaction between precipitation and hours of night ($Prec \times HoN$) as the effect of precipitation may vary with the amount of daylight.[8] We consider a specification of model (1) with a three-way interaction between cloud cover, precipitation, and hours of night as a robustness test to more flexibly allow for interaction effects between these variables.

Similarly to the specification proposed by Goetzmann et al. (2015), $\mathbf{W}_{jm}$ contains the seasonal weather averages across the years 2015-2023 for every weather variable $w \in W$ and month m for zip code j, as described in Section 2. These variables control for seasonality in both cloud cover and other weather variables.[9]

The variable $MortgageRate_m$ denotes the mortgage rate in the month m when the unit is sold. $\mathbf{X}_i$ is a vector of housing attributes for unit i to control for compositional effects, including log of age (years), log of size (square meters), an indicator taking the value of one if the housing unit is an apartment, the interaction between log of size and the apartment indicator, an indicator taking the value one if the housing unit is not a co-op unit, and an interaction between log of lot size (square meters) and a non-apartment indicator.

Weather shows a strong seasonal pattern while prices also fluctuate over time with their own seasonal trends. We expect that monthly weather controls will capture price seasonality unrelated to weather. Additionally, we include a series of fixed effects to account for any remaining confounding factors. Specifically, to control for time-invariant unobservables that may be correlated with the weather, such as institutional differences across municipalities or neighborhoods (e.g., differences in the timing of listings and sales and the composition of housing units), we include three-digit zip code fixed effects, Zip_j. To further control for local shocks at the municipality level, we include municipality-by-year-by-quarter fixed effects in the most flexible specification, μ_{it}.

To address concerns about the presence of omitted variables correlated with the weather variables or hours of night and unobserved heterogeneity in the housing units, we estimate model (1) using repeat sales with housing unit fixed effects. To mitigate concerns about housing units being sold in quick succession due to factors unobserved in the data, we also estimate the model on the repeat sales sample restricted to sales with holding time of at least 52 weeks.

[8]Because the hours of night vary seasonally by construction, controlling for the interaction between hours of night and precipitation controls for potential seasonal confounders. The Pearson correlation between the seasonal control for precipitation and interaction between $Prec$ and HoN is 0.2.

[9]Goetzmann et al. (2015) estimate a second model with de-seasoned cloud cover, where the left-hand side variable is the difference between the cloud cover and the seasonal cloud cover, i.e., $DCC = log(CC/CC^S)$. We do not estimate this model because it is equivalent to estimating the effect of deviations from the monthly zip code seasonal average while constraining the deviations to be identical across seasons, which is likely unreasonable in our setting.

3.2 Separating the effects of cloud cover on showing days versus the day of sale

We now extend model (1) to capture the effect of cloud cover on showing days and the day of sale on home prices using a two-stage approach. Because cloud cover (and precipitation to a lesser extent) is serially correlated, regressing home prices on cloud cover at both the showings and day of sale may introduce multicollinearity. To address this concern, we first residualize or orthogonalize each weather variable $w \in W$ at the day-of-sale on the four preceding days as follows:

$$w_{jt} = \alpha + \delta_1 w_{j,t-1} + \delta_2 w_{j,t-2} + \delta_3 w_{j,t-3} + \delta_4 w_{j,t-4} + \epsilon^w_{jt}, \tag{2}$$

where w_{jt} is the weather observed in zip code j on sale day t, and $w_{j,t-s}$ is the weather lagged by $s \in \{1, 2, 3, 4\}$ days.

Second, we plug in the weather residuals from model (2), $\hat{\epsilon}^w_{jt}$, which capture the day-of-sale weather effects on prices, into the following model

$$\begin{aligned} P_{it} = \alpha &+ \gamma_1 \hat{\epsilon}^{\text{CC}}_{jt} + \gamma_2 \hat{\epsilon}^{\text{CC}}_{jt} \times \hat{\epsilon}^{\text{Prec}}_{jt} + \gamma_3 \hat{\epsilon}^{\text{CC}}_{it} \times \text{HoN}_{it} \\ &+ \beta_1 \text{CC}^{MA}_{jt} + \beta_2 \text{CC}^{MA}_{jt} \times \text{Prec}^{MA}_{jt} + \beta_3 \text{CC}^{MA}_{jt} \times \text{HoN}_{jt} \\ &+ \beta_4 \text{Prec}^{MA}_{jt} + \beta_5 \text{HoN}_{jt} + \beta_6 \text{Prec}^{MA}_{jt} \times \text{HoN}_{jt} \\ &+ \beta_7 \text{PosMaxTemp}^{MA}_{jt} + \beta_8 \text{NegMaxTemp}^{MA}_{jt} + \beta_9 \text{Wind}^{MA}_{jt} \\ &+ \mathbf{W}_{jm} \beta_{\mathbf{W}} + \beta_{10} \text{MortgageRate}_t + \mathbf{X}_i \beta_{\mathbf{X}} + \mu_{it} + Zip_j + \varepsilon_{it}, \end{aligned} \tag{3}$$

The second stage estimation includes non-nested weather variables, $\hat{\epsilon}^{\text{w}}_{jt}$ for $w \in W$, ensuring that multicollinearity will not distort the estimates. We also estimate a specification of model (3) using our repeat sales sample with unit fixed effects.

4 Results

4.1 Effect of cloud cover at public showings

The effect of cloud cover at public showings are depicted in Table 2. Columns (1) to (4) show cross-sectional results, progressively adding more fixed effects, with column (3) being our preferred specification. Columns (5) and (6) show results with unit fixed effects from the repeat sales subsample. Model (1) flexibly accounts for seasonality, such that it allows to examine cloud cover effects at the seasonal level. To facilitate interpreting the cloud cover effect in the presence of interaction terms, we calculate the marginal effect of cloud cover at the means (MEM) across, and for each, season.[10] Coefficient estimates are reported in panel A and MEMs in panel B. Throughout the paper, our discussion of the effect of cloud cover will focus on the MEMs.

Estimates in column (1), featuring municipality-by-year and municipality-by-quarter fixed effects, indicate that the marginal effect of cloud cover pooled across season (Panel B) is negative but small and not significant at the 10% level. The marginal effect of cloud cover is however significant at the 1% level in the summer, suggesting that a 10% (one standard deviation from the mean) increase in cloud cover in the four days prior to the sale is associated with a 0.3% (1.4%) decrease in home prices, at the summer means of (logs of) precipitation and hours of night. This corresponds to a reduction in price of \$1092 (\$4927) for the median home. Based on the mortgage rate coefficient estimate, the cloud cover effect is equivalent to an increase in mortgage rate of 5.65% (25.49%).

Column (2) includes zip code fixed effects in addition to the fixed effects present in column (1). The pooled marginal effect of cloud cover is significant at the 5% level and twice as large (in absolute value) as that in column (1), suggesting that a 10% (one standard deviation from the mean) increase in cloud cover in the four days prior to the sale is associated with a 0.14% (0.59%) decrease in home prices. This corresponds to a reduction in price of \$490 (\$2077) for the median home, and is equivalent to an increase in mortgage rate of 1.57% (6.64%). The marginal effect of cloud cover in the spring is similar to the pooled effect because spring has similar precipitation and hours of night means as the full sample, as shown in Table 1. Thus, the pooled and spring MEMs are always quantitatively similar in all specifications. The marginal effect in the summer is twice as large (in absolute value) as the pooled effect, and is quantitatively similar to that in column (1).

Results from estimating our preferred specification with municipality-by-year-by-month and zip code fixed effects are displayed in column (3). The marginal pooled effect is similar in magnitude as that in column (2) – a 10% (one standard deviation from the mean) increase in cloud cover in the four days prior to the sale is associated with a 0.12% (0.51%) decrease in the price – although only significant at the 10% level. Spring and summer effects are similar as those in column (2), while there is no cloud cover effect in the fall.

We test the sensitivity of our results to potential strategic scheduling of showings on sunny days in column (4). Because the showings are announced in the online for-sales prospectus, typically 7 to 10 days in advance, it would be challenging to time the showings based on the weather forecast. Still, we test for this possibility. Not every day can be chosen for a showing, e.g., no Saturday. In addition, some days are more common than others depending on the local housing market (at the municipality level; see Figure A.1).

[10]We calculate the MEM using the Stata command `lincom`.

Table 2: Effect of weather on log of sell price

	Cross-section				Repeat sales	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Coefficient estimates						
Cloud cover	-0.1421***	-0.1195***	-0.1266***	-0.1358***	0.0127	0.0168
	(0.0469)	(0.0355)	(0.0366)	(0.0365)	(0.0506)	(0.0499)
Prec.	0.0180**	0.0104	0.0110	0.0114	0.0030	-0.0008
	(0.0091)	(0.0069)	(0.0070)	(0.0070)	(0.0098)	(0.0094)
Cloud cover x Prec.	-0.0198**	-0.0049	0.0001	-0.0014	0.0069	0.0028
	(0.0084)	(0.0064)	(0.0064)	(0.0064)	(0.0087)	(0.0086)
Cloud cover x HoN	0.0684***	0.0491***	0.0510***	0.0561***	-0.0173	-0.0156
	(0.0205)	(0.0156)	(0.0160)	(0.0160)	(0.0223)	(0.0220)
Prec. x HoN	-0.0046	-0.0032	-0.0040	-0.0039	-0.0012	0.0012
	(0.0038)	(0.0029)	(0.0029)	(0.0029)	(0.0041)	(0.0040)
Cloud cover (season)	4.3121***	0.6147***	0.5576***	0.5651***	0.0568	0.0713
	(0.0816)	(0.0725)	(0.0724)	(0.0723)	(0.1029)	(0.1039)
Prec. (season)	-0.4681***	-0.0854***	-0.0888***	-0.0897***	-0.0212	-0.0173
	(0.0094)	(0.0089)	(0.0089)	(0.0089)	(0.0136)	(0.0135)
Hours of night	0.1498***	-0.0194**	-0.0229**	-0.0278***	0.0145	0.0097
	(0.0119)	(0.0096)	(0.0098)	(0.0098)	(0.0139)	(0.0135)
Mortgage rate	-0.0552***	-0.0893***	-0.1116***	-0.0988**	-0.1264**	-0.1296**
	(0.0155)	(0.0119)	(0.0390)	(0.0390)	(0.0553)	(0.0531)
N	52,911	52,909	52,906	52,906	12,170	11,412
Adj. R. sq.	0.7781	0.8709	0.8728	0.8734	0.9593	0.9621
Fixed effects	MY & MQ	MY & MQ	MYQ	MYQ & MDW	YQ	YQ
Zip code FE		✓	✓	✓		
Unit FE					✓	✓
Holding time (weeks)						≥52
Panel B: Marginal effects at the means, pooled and by season						
Pooled	-0.0073	-0.0140**	-0.0121*	-0.0112*	-0.0196**	-0.0154*
	(0.0083)	(0.0063)	(0.0064)	(0.0064)	(0.0090)	(0.0088)
Spring	-0.0071	-0.0151**	-0.0140**	-0.0131**	-0.0199**	-0.0152*
	(0.0079)	(0.0060)	(0.0061)	(0.0061)	(0.0086)	(0.0085)
Summer	-0.0312***	-0.0305***	-0.0288***	-0.0297***	-0.0134	-0.0101
	(0.0110)	(0.0083)	(0.0085)	(0.0085)	(0.0117)	(0.0114)
Fall	0.0163	0.0039	0.0071	0.0098	-0.0253**	-0.0210*
	(0.0115)	(0.0088)	(0.0090)	(0.0090)	(0.0126)	(0.0125)

Notes: The table shows results from regressing the log of home prices on cloud cover and other controls. All variables are natural logarithms. Weather variables are measured as the moving averages of the four days prior to the sale. HoN denotes the hours of night. Fixed effects are abbreviated as follows: M is municipality, Y is year, Q is quarter, and DW is day-of-the-week. Thus, MYQ denotes the municipality-by-year-by-quarter fixed effects. All regressions include the weather controls maximum temperature and wind speed and their seasonal monthly averages. All regressions control for age and mortgage rate. Cross-section regressions control for attributes: size (square meters), an indicator taking the value of one if the housing unit is an apartment, the interaction between size and the apartment indicator, an indicator taking the value of one if the housing unit is a non-co-op unit, and an interaction between lot size (square meters) and a non-apartment indicator. Heteroskedasticity-robust standard errors are given in parenthesis. * $p<0.1$, ** $p<0.05$, *** $p<0.01$.

Therefore, we include day-of-the-week (for the sale date; Monday-Sunday) by municipality fixed effects to account for institutional differences across municipalities. Estimates in column (4) are quantitatively similar to those in column (3), suggesting that variation in the day-of-the-week sale is not a concern.

The cross-sectional estimates (columns (1)-(4)) may suffer from unobserved heterogeneity in the housing units that is correlated with cloud cover and home prices. For instance, the value of a balcony likely depends on its orientation and sunlight exposure, and architectural choices such as window size and floor plan design can influence the amount of sunlight that enters the housing unit. This motivates us to estimate model (1) with unit fixed effects in columns (5) and (6). Column (5) uses the unrestricted repeat sales sample. Estimates indicate that a 10% (one standard deviation from the mean) increase in cloud cover in the four days prior to the sale is associated with a 0.2% (0.83%) decrease in home prices, significant at the 5% level, which corresponds to a reduction in price of $686 ($2908) for the median home. This is equivalent to an increase in mortgage rate of 1.55% (6.57%). The key difference with the cross-section analysis is that the effect in the summer is more than halved compared to column (3) and is no longer significant at the 10% level, while the effect in the fall is now negative and significant, with a MEM of -0.0253.[11] These results suggest that during periods of the year with longer and brighter days, sunlight at the showings does not have an effect on prices, but when the days are darker, the effect of sunlight becomes more appreciated. This result is notable as it highlights how buyers' price formation are influenced by the level of sunlight during showings, conditional on the length of daylight hours.

Column (6) restricts the sample to repeat sales at least 52 weeks apart. Estimates for both repeat sales samples are qualitatively similar, although slightly smaller (in absolute value) and less significant in the restricted sample (column (6)).

4.1.1 Robustness and sensitivity

To investigate the robustness of our results, we run a battery of tests in Appendix B.1. Specifically, we re-estimate model (1) with alternative controls, including the variable proposed by Kamstra et al. (2003) to account for seasonal affective disorder (SAD), alternative fixed effects, and with three-way interaction specifications. The results are in general robust to these alternative choices. We further run sensitivity tests around our choice of the four day moving average window, between two and six days – we do not in-

[11]Controlling for time-invariant unobserved heterogeneity results in a switch in the sign and magnitude of the coefficient estimates, specifically for cloud cover and its interaction with hours of night (Panel A). Even though the repeat sales point estimates are no longer significant, it is clear that the relevant variable is the interaction, which ultimately makes the effect at the sample means significant, except in the summer.

crease the moving average window beyond six days because it would then start resembling the seasonal controls. The four day moving average window is empirically and institutionally motivated. Again, the results are overall robust to such alternative choices. Last, we examine the sensitivity of our results to restrictions of the time-on-market (TOM). Restricting TOM in the sample increases the likelihood that public showings take place closer to the sale and, thus, during the four-day moving average. The maximum TOM in the sample is 22 days. We run sensitivity tests by restricting the TOM down to 10 days. Results indicate that the shorter the TOM, the stronger the cloud cover effect becomes (up to twice as large in magnitude). This is consistent with the fact that our four day moving average is more likely to measure the cloud cover at the showings.

4.1.2 Validity

To test the validity of our main results, we randomly draw, without replacement, weather realizations from the full time series of observations in the time period 2015-2023. We use two approaches for assigning random draws of weather to every transaction in our sample. We either draw a) one variable at a time, or b) the full realization of weather. The procedure is repeated 1,000 times. We report the mean and standard deviation of coefficient estimates. We also calculate the MEMs at each iteration and report the mean and standard deviation of the MEM estimates. Results for both the cross-sectional and repeat sales analyses (Tables B.12 and B.13, respectively) indicate that the MEMs of cloud cover are close to zero.

4.1.3 Results' heterogeneity across subsamples

We investigate how the effects of cloud cover at the showings on home prices differ across unit types, i.e., apartments and non-apartments, and across buyer characteristics. Table B.5 indicates an effect of cloud cover on home prices for apartments similar to that found in Table 2 for the repeat sales analysis, while the effect is insignificant for non-apartments. These findings may be explained by the fact that non-apartments are likely more heterogeneous than apartments, such that the potential for cloud cover capturing unobserved time-varying heterogeneity is more pronounced in the case of non-apartments. Furthermore, the absence of cloud cover effect for non-apartments is also suggesting that cloud cover, or its counterpart sunlight, does not dramatically affect prospective buyers' valuation of a home's attributes, e.g., a single family home's curb appeal or garden appearing more attractive when the sun shines. This sunshine-enhanced effect of a home's features would be arguably more pronounced for non-apartments than for apartments, for which outdoor features are often less salient.

Using information on buyer characteristics at the end of the year preceding the home purchase, we now investigate how different groups of buyers may be susceptible to cloud cover. Due to concerns about omitted variables in the cross-sectional specification, we segment the repeat sales sample with unit fixed effects based on buyer characteristics in the last sales within a repeat sales sequence. The repeat sales results, shown in Figure 2, suggest that in general transactions with a single buyer are affected by sunlight, while there is no effect with multiple buyers. Similarly, transactions with a single adult in the household seem affected by sunlight, while transactions with buyers belonging to two-adult households are not affected. The effect is of comparable magnitude, although less significant, for households without children, while those with children are not affected. Furthermore, the effect of cloud cover is more pronounced for buyers with total gross household wealth below the sample median than for those with total gross household wealth above the sample median. Last, men buying alone do not appear to be affected by cloud cover, while women buying alone do,[12] and first-time buyers seem more affected than non-first timer buyers. Taken together, these results provide support to the theory that buyers with less bargaining power or purchasing power are more likely to be affected by cloud cover-induced mood.[13] (Cross-sectional results are reported in Tables B.6, B.7, and B.8.)

4.2 Effect from the day of the sale

We estimate the marginal effect of cloud cover at the showings (Panel B) and at the day of sale (Panel A) in Table 3, in which the first four columns use the cross-section sample and the last four columns use the repeat sales sample.

First, we estimate the effect of cloud cover at the day of the sale using model (1) where we simply replace the weather for the four day moving average prior to the sale by the weather at the sale date (columns (1) and (5)). Second, we use both the four day moving average and day of sale weather in model (1) (columns (2) and (6)). Third, we estimate the effect of cloud cover at the day of the sale by first residualizing it using model (2), and then estimate model (3) (columns (3) and (7)). Fourth, we jointly estimate the effect of cloud cover at the day of sale by first residualizing it using model (2) and at weather at the showings, and then estimate model (3) (columns (4) and (8)). Note that because the preferred model (model (3)) uses weather on the day of the sale consisting of *unexplained variation* after accounting for four lags, these *residuals* are simply shifted versions of the

[12]Closer examination reveals that single men are affected to the same extent as single women when restricting the analysis to single men buying from non-single men versus single women buying from non-single women. The non-significant result for single men could be due to specific unobserved heterogeneity in housing units transacted between single men. Results are available upon request.

[13]The results for households without children are consistent with this theory because they are more likely to be single-adult households and single buyers. In turn, because gross wealth is calculated at the household level, single-adult households are more likely to be in the lower half of the wealth distribution.

Figure 2: Buyer heterogeneity - repeat sales

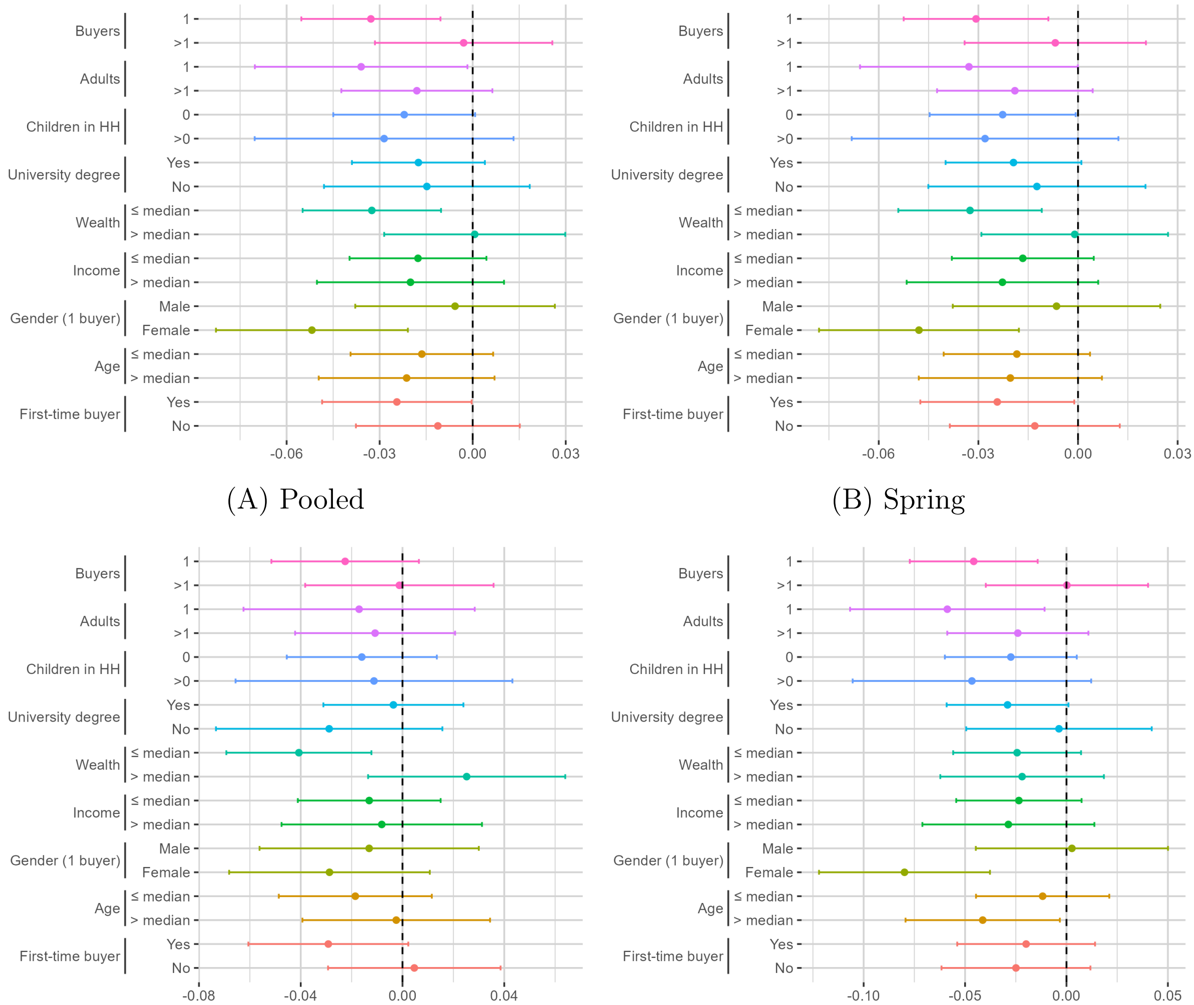

(A) Pooled

(B) Spring

(C) Summer

(D) Fall

Notes: The figures show marginal effects at the means from regressing log of the selling price on the cloud cover at the showings, in which the selling price includes common debt if any. The effects are evaluated at the sample means of the full sample, as well as by season, without conditioning on segmentation. We also report the 95% confidence intervals for the MEMs, calculated using heteroskedasticity-robust standard errors. The samples are segmented by the characteristics of the buyers in the final sale of the repeat sales sequence. *Buyers* denotes segmenting by the number of buyers. *Adults in HH* denotes segmenting by the number of adults in the household, focusing on transactions in which all buyers belong to the same household. *Children in HH* denotes segmenting by the number of children in the household, defined as those up to 17 years old, focusing on transactions in which all buyers belong to the same household. *University degree* denotes segmenting by whether any of the buyer have a university degree. *Wealth* denotes segmenting by the total gross wealth among the buyers. *Income* denotes segmenting by the total gross income among the buyers. Wealth and income are deflated to 2015 price levels using the consumer price index. *Gender* denotes segmenting by the gender of buyers, focusing on transactions with one buyer. *Age* denotes segmenting by the average age among the buyers. *First-time buyers* denotes segmenting by whether none of the buyers already own a housing unit, based on the taxable value of their primary home. These results are also provided in Tables B.9, B.10 and B.11.

actual weather. As a result, the magnitude of the estimated coefficients is comparable across specifications in Table 3 and between days-of-the-sale and showings.

Across all specifications, the marginal effect of cloud cover at the day of sale on home prices appears both economically very small and insignificant at the 10% level. The marginal effects of cloud cover at the showings on home prices are qualitatively similar to those in Table 2, although slightly more pronounced and significant, in particular when using the unrestricted repeat sales sample. (Results for the repeat sales sample with the 52-week minimum holding time restriction are qualitatively similar, although also somewhat less pronounced; see Table B.4.)

4.3 Effect of cloud cover on market liquidity

Taken together, the results on the effect of cloud cover at the showings and at the day of sale on prices suggest the housing market is affected by sunlight at the showings but not at the day of the sale. One mechanism might be that cloud cover-induced mood dampens buyer interest leading to reduced market liquidity.[14] An alternative mechanism – with the opposite effect on market liquidity – might be that cloud cover reduces buyers' desire for sunshine-dependent activities, thereby lowering their opportunity costs of visiting a home and making attending showings more attractive. We investigate in Table 4 how cloud cover might affect market liquidity using three different measures: (1) the number of showing participants, (2) the number of bidders, and (3) the sell-ask spread. We estimate model (1) with each of our three liquidity measures as the dependent variable.[15] The marginal effects do not indicate that the number of showing participants, number of bidders, or the sell-ask spread are affected by cloud cover.

4.4 Discussion

What mechanisms might explain how sunlight affects home prices? One possible explanation is that sunlight influences buyer interest by increasing the opportunity cost of attending showings – e.g., buyers may prefer to engage in outdoor activities when the sun shines. However, the results in Section 4.3 indicate that this mechanism is unlikely, as cloud cover does not correlate with changes in number of prospective buyers.

Another possible mechanism, discussed by Hirshleifer and Shumway (2003), is the misattribution of mood as information cue. Our findings are consistent with this mechanism

[14]Because the decision to list a unit takes place several weeks before the listing, we assume that sellers' decision to put their home on the market is not affected.

[15]The number of showing participants and number of bidders analyses use two subsamples of the supplementary data described in Section 2, from October 2021 to June 2023 and August 2018 to June 2023, respectively. Because those two samples cover relatively short time periods, we do not estimate a unit fixed effect specification.

Table 3: Marginal effects of cloud cover at showings and day of sale on log of home prices (evaluated at means; MEMs)

	Cross-section				Repeat sales			
	Non-residualized		Residualized		Non-residualized		Residualized	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Day of the sale								
All	0.0021	0.0026	0.0016	0.0005	-0.0044	-0.0033	0.0014	-0.0018
	(0.0041)	(0.0042)	(0.0036)	(0.0036)	(0.0060)	(0.0061)	(0.0053)	(0.0054)
Spring	0.0001	0.0007	0.0006	-0.0007	-0.0039	-0.0030	0.0005	-0.0024
	(0.0037)	(0.0037)	(0.0034)	(0.0035)	(0.0054)	(0.0055)	(0.0050)	(0.0051)
Summer	-0.0011	0.0003	0.0005	-0.0019	-0.0054	-0.0062	-0.0014	-0.0045
	(0.0050)	(0.0051)	(0.0046)	(0.0048)	(0.0074)	(0.0075)	(0.0069)	(0.0070)
Fall	0.0077	0.0073	0.0040	0.0044	-0.0041	-0.0007	0.0052	0.0017
	(0.0061)	(0.0062)	(0.0053)	(0.0054)	(0.0089)	(0.0091)	(0.0078)	(0.0079)
Panel B: Days prior to the sale, MA(4)								
Pooled		-0.0117*		-0.0116*		-0.0195**		-0.0200**
		(0.0065)		(0.0065)		(0.0091)		(0.0091)
Spring		-0.0135**		-0.0137**		-0.0194**		-0.0199**
		(0.0062)		(0.0062)		(0.0088)		(0.0088)
Summer		-0.0283***		-0.0297***		-0.0111		-0.0127
		(0.0085)		(0.0086)		(0.0118)		(0.0118)
Fall		0.0071		0.0090		-0.0281**		-0.0274**
		(0.0092)		(0.0092)		(0.0129)		(0.0129)
N	52,906	52,906	52,906	52,906	11,706	11,706	11,706	11,706
Adj. R. sq.	0.8728	0.8728	0.8728	0.8728	0.9598	0.9598	0.9598	0.9598
Fixed effects	MYQ	MYQ	MYQ	MYQ	YQ	YQ	YQ	YQ
Zip code FE	✓	✓	✓	✓				
Unit FE					✓	✓	✓	✓

Notes: The table shows marginal effects at the means from regressing the cloud cover on log of the selling price, in which the selling price includes common debt if any. All results are from estimating specification (3). Cross-sectional results are in columns (1)-(4) and unit fixed effect are in columns (5)-(8). Column (1) estimates the model with the weather at the day of the sale, while column (2) adds the moving averages as well. Columns (3) and (4) substitute the raw weather at the day of the sale with residualized weather. The same goes for columns (5)-(8). MYQ and YQ denotes the municipality-by-year-by-quarter and year-by-quarter fixed effects, respectively. All regressions include the weather controls maximum temperature and wind speed and their seasonal monthly averages. Heteroskedasticity-robust standard errors are given in parenthesis. Significance: * p<0.1, ** p<0.05, *** p<0.01.

as we observe that sunlight influences prices during the initial information processing stage – specifically at the showings – and this effect is more pronounced during periods with shorter, darker days.[16] Thus, sunlight seems to impact buyers' mood most when it is generally in shorter supply.

Psychology studies point to several cognitive processes that may explain how mood

[16] Sunlight may also act as a mood primer before showings, influencing buyer mood on their way to the property, such as during their travel to the viewing.

Table 4: Marginal effects of cloud cover on the number of showing participants, the bidders, and the sell-ask spread (evaluated at means; MEMs)

	Showing participants		Bidders		Sell-Ask	
	(1)	(2)	(3)	(4)	(5)	(6)
Marginal effects at the means, pooled and by season						
Pooled	0.0087	-0.0142	-0.0109	-0.0118	-0.0004	-0.0061
	(0.0408)	(0.0409)	(0.0182)	(0.0183)	(0.0026)	(0.0077)
Spring	0.0061	-0.0144	-0.0081	-0.0102	-0.0018	-0.0062
	(0.0388)	(0.0389)	(0.0169)	(0.0170)	(0.0025)	(0.0074)
Summer	0.0332	-0.0055	-0.0079	-0.0161	-0.0044	0.0029
	(0.0540)	(0.0544)	(0.0232)	(0.0233)	(0.0035)	(0.0096)
Fall	-0.0121	-0.0229	-0.0164	-0.0097	0.0053	-0.0151
	(0.0654)	(0.0651)	(0.0277)	(0.0278)	(0.0036)	(0.0106)
N	10,678	10,672	26,262	26,260	52,906	12,170
Adj. R. sq.	0.1521	0.1813	0.0953	0.0995	0.1999	0.2156
Fixed effects	MYQ(s)	MYQ(s)	MYQ	MYQ	MYQ	YQ
County x DW		✓(s)		✓		
Zip code FE	✓	✓	✓	✓	✓	
Unit FE						✓

Notes: The table shows marginal effects at the means from regressing log of the number of showing participants on the cloud cover at showings, the log of the number of bidders on the cloud cover at the day of the sale, and the sell-ask spread on the cloud cover the four days prior to the sale. For the showing participation results, the sample period is October 2021 to June 2023, and each row in the sample is a unique showing date-transaction instance. If there is more than one showing on a particular day, we use the sum of the showing participants on that day as the participation number. For the bidders results, sample period is August 2018 to June 2023. The sell-ask spread is defined as (sell-ask)/ask, and the sell-ask model estimated using the main sample with the same explanatory variables and controls. All variables are in natural logarithms except the sell-ask spread. MYQ denotes the municipality-by-year-by-quarter fixed effects, and DW is the day-of-the-week. Showing participation fixed effects are denoted with (s) because these are for the showing dates and not the sales date. All regressions include the weather controls maximum temperature and wind speed and their seasonal monthly averages. All regressions control for age and mortgage rate, size (square meters), an indicator taking the value of one if the housing unit is an apartment, the interaction between size and the apartment indicator, an indicator taking the value of one if the housing unit is an owner-occupier unit, and an interaction between lot size (square meters) and a non-apartment indicator. Heteroskedasticity-robust standard errors are given in parenthesis. Significance: * $p<0.1$, ** $p<0.05$, *** $p<0.01$.

gets misattributed as information cue: (1) mood-tainted memory encoding and (2) congruent mood recall. First, when visiting housing units, buyers may encode their memories with their current mood (e.g., Bower, 1981; Bower and Forgas, 2000; Forgas, 2017). When recalling the specifics about the housing unit at a later time, these memories will retain the "emotional tone" present at the time the memories were encoded. Thus, when recalling the memories from the showings, the emotional tone of the memories may bias the decision making on the day of the sale. Second, the current mood may amplify the effect through mood-congruent recall, in which the tone of the memories that are congruent with the current mood are more available to the buyer (Isen et al., 1978; Bower, 1981). For example, when the sun shines, it may be easier for the prospective buyer to recall attributes of the housing unit that she remembers positively. To investigate whether sunlight on the day of the sale affects prices through a mood-congruent recall channel, we re-estimate model (1) excluding observations with dissimilar cloud cover at the showings and the sale. The results, presented in Table B.14, show a slightly stronger effect when dropping dissimilar cloud cover, consistent with the mood-congruent recall channel. Furthermore, decision processes that require much effort, cognitive attention, and are stressful, as buying a home arguably is, are more prone to being affected by mood (Forgas, 1995). We find that single buyers, first-time buyers, buyers from one-adult households, and buyers with lower wealth appear more susceptible to cloud cover. These findings are consistent with the theory that having less bargaining power or a weaker position in the housing market (in terms of experience in the housing market and wealth) may make such buyers more prone to being affected by mood.

Last, another possible mechanism is that sunlight enhances the perceived value of a home's features or curb appeal, e.g., a terrace or garden may look more appealing in bright sunshine. Our results suggest no effect of cloud cover for non-apartments, e.g., single family homes with a garden, where curb appeal would be most likely to benefit from sunshine. Instead, we find evidence that cloud cover reduces apartment prices. Taken together, these findings are consistent with the mood channel mechanism – although we cannot rule out that sunshine directly affects a prospective buyer's valuation of a home's features.

5 Conclusions

Cloud cover has been shown to negatively correlate with several economic and financial outcomes. For example, Goetzmann et al. (2015) demonstrate that cloud cover directly increases institutional investors' pessimism and perception that stocks are overpriced. This paper provides evidence that cloud cover-induced mood affects a new asset class, the residential real estate market, and uncovers the channels through which it influences price formation.

Using data from the Norwegian housing market, for which we are able to relatively precisely identify the different stages of the sales process, including showing days and offer submission, we investigate whether home buyers are sensitive to cloud cover during the sales process. Despite housing being often characterized as the largest and most significant asset for households, we do find evidence that home prices are negatively affected by cloud cover – a standard deviation increase in cloud cover from the mean at the showings reduces home prices by 0.83% (p-value<0.05). For the median home price in our sample, this corresponds to a reduction of $2,908 in price and is equivalent to a 6.57% increase in the mortgage rate. Results are robust to a battery of robustness tests and alternative specifications. Furthermore, we find suggestive evidence that apartment buyers, first-time buyers, and less wealthier buyers tend to be most influenced by mood changes induced by cloud cover. Last, we demonstrate that cloud cover primarily influences price formation during the public showings of homes for sale. In contrast, we find no evidence that cloud cover affects the number of prospective buyers attending showings, the price formation during the auctions or the number of bidders in auctions. Our findings suggest that sunshine or cloud cover, through their effect on mood, may bias residential real estate prices, with implications for household savings and debt, and the macroeconomy.

Future research could elicit via a survey prospective buyers' housing market perceptions at different points in time and/or location as in Goetzmann et al. (2015) to establish which connections prevail between sunshine or cloud cover and beliefs about a home's value – one notable distinction between real estate and stock market indices is the existence of private values.

References

Bodoh-Creed, A. L. (2020). Mood, memory, and the evaluation of asset prices. *Review of Finance*, 24(1):227–262.

Bower, G. H. (1981). Mood and memory. *American psychologist*, 36(2):129.

Bower, G. H. and Forgas, J. P. (2000). Affect, memory, and social cognition. In Eich, E., Kihlstrom, J. F., Bower, G. H., Forgas, J. P., and Niedenthal, P. M., editors, *Cognition and Emotion*. Oxford University Press.

Campbell, J. Y. and Cocco, J. F. (2007). How do house prices affect consumption? evidence from micro data. *Journal of monetary Economics*, 54(3):591–621.

Clore, G., Schwarz, N., and Conway, M. (1994). Affective causes and consequences of social information processing. In Wyer Jr, R. S. and Srull, T. K., editors, *Handbook of Social Cognition: Volume 1: Basic Processes*, volume 1, pages 323–417. Psychology Press.

Copernicus Climate Change Service (2024a). Cloud properties global gridded monthly and daily data from 1982 to present derived from satellite observations. Copernicus Climate Change Service (C3S) Climate Data Store (CDS).

Copernicus Climate Change Service (2024b). ERA5-Land hourly data from 1950 to present. Copernicus Climate Change Service (C3S) Climate Data Store (CDS).

Copernicus Climate Change Service (2024c). Nordic gridded temperature and precipitation data from 1971 to present derived from in-situ observations. Copernicus Climate Change Service (C3S) Climate Data Store (CDS).

Cunningham, M. R. (1979). Weather, mood, and helping behavior: Quasi experiments with the sunshine samaritan. *Journal of personality and social psychology*, 37(11):1947.

Denissen, J. J., Butalid, L., Penke, L., and Van Aken, M. A. (2008). The effects of weather on daily mood: a multilevel approach. *Emotion*, 8(5):662.

Forgas, J. P. (1995). Mood and judgment: the affect infusion model (aim). *Psychological bulletin*, 117(1):39.

Forgas, J. P. (2017). Mood effects on cognition: Affective influences on the content and process of information processing and behavior. In Jeon, M., editor, *Emotions and Affect in Human Factors and Human-Computer Interaction*, pages 89–122. Elsevier.

Goetzmann, W. N., Kim, D., Kumar, A., and Wang, Q. (2015). Weather-induced mood, institutional investors, and stock returns. *The Review of Financial Studies*, 28(1):73–111.

Goetzmann, W. N. and Zhu, N. (2005). Rain or shine: where is the weather effect? *European Financial Management*, 11(5):559–578.

Gourley, P. (2021). Curb appeal: how temporary weather patterns affect house prices. *The Annals of Regional Science*, 67(1):107–129.

Hirshleifer, D. and Shumway, T. (2003). Good day sunshine: Stock returns and the weather. *The journal of Finance*, 58(3):1009–1032.

Isen, A. M., Shalker, T. E., Clark, M., and Karp, L. (1978). Affect, accessibility of material in memory, and behavior: A cognitive loop? *Journal of personality and social psychology*, 36(1):1.

Jacobsen, B. and Marquering, W. (2008). Is it the weather? *Journal of Banking & Finance*, 32(4):526–540.

Johnson, E. J. and Tversky, A. (1983). Affect, generalization, and the perception of risk. *Journal of personality and social psychology*, 45(1):20.

Kamstra, M. J., Kramer, L. A., and Levi, M. D. (2003). Winter blues: A SAD stock market cycle. *American Economic Review*, 93(1):324–343.

Kamstra, M. J., Kramer, L. A., Levi, M. D., and Wang, T. (2014). Seasonally varying preferences: Theoretical foundations for an empirical regularity. *The Review of Asset Pricing Studies*, 4(1):39–77.

Karlsson, K.-G., Riihelä, A., Trentmann, J., Stengel, M., Solodovnik, I., Meirink, J. F., Devasthale, A., Jääskeläinen, E., Kallio-Myers, V., Eliasson, S., Benas, N., Johansson, E., Stein, D., Finkensieper, S., Håkansson, N., Akkermans, T., Clerbaux, N., Selbach, N., Marc, S., and Hollmann, R. (2023). CLARA-A3: CM SAF cLoud, Albedo and surface RAdiation dataset from AVHRR data - Edition 3.

Kelly, P. J. and Meschke, F. (2010). Sentiment and stock returns: The sad anomaly revisited. *Journal of Banking & Finance*, 34(6):1308–1326.

Leamer, E. E. (2007). Housing is the business cycle. Working Paper 13428, National Bureau of Economic Research.

Loewenstein, G. (2000). Emotions in economic theory and economic behavior. *American Economic Review*, 90(2):426–432.

Muñoz Sabater, J. (2019). ERA5-Land hourly data from 1950 to present. Copernicus Climate Change Service (C3S) Climate Data Store (CDS).

Saunders, E. M. (1993). Stock prices and wall street weather. *The American Economic Review*, 83(5):1337–1345.

Schwarz, N. (2011). Feelings-as-information theory. In Van Lange, P. A. M., Higgins, E. T., and Kruglanski, A. W., editors, *Handbook of Theories of Social Psychology*, volume 1, pages 289–308. Sage.

Schwarz, N. and Clore, G. L. (1983). Mood, misattribution, and judgments of well-being: Informative and directive functions of affective states. *Journal of personality and social psychology*, 45(3):513.

Schwarz, N. and Clore, G. L. (2007). Feelings and phenomenal experiences. In Kruglanski, A. W. and Higgins, E. T., editors, *Social psychology: Handbook of basic principles*, pages 385–407. Guilford, New York, NY, 2 edition.

Statistics Norway (2024). Alders- og kjønnsfordeling i kommuner, fylker og hele landets befolkning (k) 1986 - 2024.

Tveito, O., Førland, E., Heino, R., Hanssen-Bauer, I., Alexandersson, H., Dahlström, B., Drebs, A., Kern-Hansen, C., Jónsson, T., E., V.-L., and Westman, Y. (2000). Nordic Temperature Maps. Technical Report 9/00 KLIMA, Norwegian Meteorological Institute.

Tveito, O. E., Bjørdal, I., Skjelvåg, A. O., and Aune, B. (2005). A gis-based agro-ecological decision system based on gridded climatology. *Meteorological Applications*, 12(1):57–68.

A More on the data

This section gives more information about the data we use in the analysis, including the data cleaning process and more statistics.

A.1 Housing transactions

Our main data set of transactions from Eiendomsverdi AS is cleaned as follows. We start by dropping transactions with a larger TOM than two times the median, meaning, we drop units that take longer than 22 days to sell. Then we trim on the 1st and 99th percentiles of ask and sell price by municipality and year. We also winsorize the living area size, by municipality and year, on the 1st and 99th percentiles. Finally, we restrict to the seasons spring, summer, and fall, and keep the south-eastern municipalities. These include Asker, Bærum, Drammen, Fredrikstad, Kristiansand, Lillestrøm, Moss, Nordre Follo, Oslo, Sandefjord, Sarpsborg, Skien, and Tønsberg. Based on numbers from the Norwegian statistics agency (Statistics Norway), in 2021 these municipalities populated about 30 percent of the total population in Norway, meaning about 1,650,000 out of 5,400,000 (Statistics Norway, 2024). The data cleaning process is presented in Table A.1.

These data are merged with buyer characteristics for the purpose of investigating whether different buyer groups are more susceptible to being influenced by cloud cover. Buyer characteristics from the end of the year before the sale are used to account for their situation at the time of purchase. Note that the buyer characteristics dataset ends in 2019. Therefore, when investigating buyer heterogeneity in Subsection 4.1.3, transactions in 2021 use buyer characteristics from the end of 2019. Summary statistics for the buyer characteristics are presented in Table A.3.

A.2 Weather data

The weather data from the Copernicus Climate Change Service are collected for the period 2015-2023.

Cloud cover: We use cloud cover measured as a daily daytime average. Daytime averages are not available in darker periods, starting in mid October and ending in February. When the daytime cloud cover average is not available for a particular day, we use the daily cloud cover average. The cloud cover median grid cell size is 336 square kilometers with rectangular grid cells. Cloud cover is provided on a scale from 0 to 100, in which 0 is clear sky and 100 is fully cloud covered sky. We scale this to being between 0 and 1.

Precipitation and maximum temperature: The consolidated data with precipitation and maximum temperatures are provided on a one-by-one kilometer resolution. Precipitation is provided in millimeters per day and is accumulated from 6:00 AM UTC the reported date to 6:00 AM UTC the day after. The maximum temperature is provided in Kelvin which we convert to Celsius by subtracting 273.15. It is the maximum in the time interval between 6:00 PM UTC the date before and 6:00 PM UTC the date of the reported observation. Most likely the temperature we use is the one falling on the day of the reported observation.

Wind speed: Wind speed is a short forecast based on the ERA5 model, and we find that the median grid cell size is approximately 100 square kilometers with rectangular grid cells. Wind speed is found from the (10m) u- and v-components of wind using the Pythagorean Theorem: $\mathrm{WS} = \sqrt{u^2 + v^2}$. The forecasts we use are made for 12:00 PM UTC, and it is meant as a control variable because the local wind speed is highly dependent on local factors such as topology and building structure. Note that Norway is located in the CET timezone and uses daylight saving times, so that during winter time the wind speed is for 1:00 PM and during summer time it is for 2:00 PM.

Aggregation to the zip code level: The weather data are converted from grids to zip codes, taking the average of the grid values if more than one grid is overlapping the zip codes. Distributions of the four weather variables are provided in Table A.4 and Figure A.4.

Hours of night: We calculate hours of night (hereafter HoN) using the the method proposed by Kamstra et al. (2003) to calculate their SAD variable, although we do not truncate the hours of night to zero in the spring and summer as they do for SAD, and do not subtract 12.

Log transformation of weather variables: All weather variables are log transformed plus one due to the presence of zeros. The maximum temperatures are provided in Celsius, therefore, we take the natural logarithm of the absolute value of these plus one, then create two temperature variables: one for positive temperatures and one for negative temperatures.

A.3 Supplementary data with showing participation and bidders

The supplementary data set is smaller than our main sample, and contain observations from both urban and rural areas. The subsampling includes keeping units sold in municipalities below latitude 61.5 and then dropping those located in the old counties (fylker) Hordaland, Rogaland, and Vest-Agder, except for the muncipalities Kristiansand and Vennesla. Kristiansand is in our main sample, and is located in the south-west of our sample region. We include rural areas because of the limited sample size. In this data set, we are not able to differentiate between type of buyers, that being whether they are people, firms or official institutions. We expect the sample to be prone to much more unobserved heterogeneity because of the inclusion of rural areas and buyer types.

Therefore, we restrain to not replicate the sell price estimates using these data. In total, there are 104 municipalities in this sample. Although we keep all municipalities in the south-east of Norway to match our main study area, the DNB Eiendom data have much fewer observations than our main sales transaction sample. We trim by county instead of municipalities. The data cleaning process is presented in Table A.2.

The cleaned supplementary data set is used to motivate the choice of the four day moving average to measure the cloud cover at the showings. In addition to the information in Subsection 2.2 that motivates the moving average window, there are some relevant nuances. When there is single day for the public showing(s), meaning that all showings are conducted during a single day, 78% of sales happen one day later, and 91% are sold within four days. When showings take place over several days, 73% of sales happen four days after the first showing. Further up in these distributions the days between showings and sales increase rapidly. The reason for this is that some housing units are not sold after the first auction, leading to new rounds of showings. We cannot exclude the possibility that prospective buyers visiting the unit in the first round of showings will not buy the unit a couple of weeks later at a second auction. We assume that such prospective buyers would either have lost interest in the unit or would attend the new round of showing(s) and that this latest round of showings would be most relevant for their decision-making. Given this assumption, the relevant showings are those that take place closest to the sale date.

A.4 Tables and figures

Table A.1: Cleaning steps, main data

Cleaning step	# of transactions
Initial	160,541
Keep relevant areas	133,749
Keep sales with TOM of at most 2 times the median (22 days)	97,666
Keep sales with TOM of at least 7 days	84,974
Trimming	82,800
Drop the winter	68,926
Keep south-east municipalities	56,226

Notes: The table shows the steps taken for cleaning the main data set and how many transactions remain after each step.

Table A.2: Cleaning steps, data with observed showing dates

Cleaning step	# of transactions
Initial	95,387
Keep on-market sales	94,923
Keep second-hand market realtor offices	94,655
Drop duplicates	94,655
Drop sales with more than one accepted bid	93,679
Keep sales between 2018 and 2023	92,115
Drop bids with missing timestamp for bids	92,084
Drop sales without sales date	90,801
Keep municipalities below latitude 61.5	79,044
Keep auctions lasting at most 7 days (received first bid to sale)	66,916
Keep sales with showings after the listing (dropping previously withdrawn units)	52,988
Keep sales with TOM of at most 22 days (same as used on the main data set)	43,932
Drop units with first (and last) showing after the sale	42,789
Drop sales with missing size	42,761
Drop sales assigned to the wrong county (fylke)	42,748
Trimming	41,646
Drop the winter	35,077
Keep south-east municipalities	29,372
Keep transactions with participation at precise showing dates	10,209

Notes: The table shows the steps taken for cleaning the data set with the showing dates and how many transactions remain after each step. Note that the last cleaning step is used for the showing participation regressions in Table 4, while the regressions for the number of bidders use the sample from the second to last step. Also note that this data set is initially at the bid level, and to keep well-behaved sales we need to include some additional restrictions. These data include typos and other clerical errors which we sort out. Trimming is done as for our main data set, but we do the trimming by county (fylke) instead of at the municipality level.

Table A.3: Summary statistics - buyers

	Mean	SD	25th pct.	Median	75th pct.
Panel A: All sales (N = 55,823)					
# of buyers	1.51	0.52	1.00	2.00	2.00
All buyers from same HH	0.84				
Adults in HH	1.88	0.84	1.00	2.00	2.00
Any children in HH	0.32				
Any buyers with university degree	0.29				
Gross wealth (in 1,000 USD)	225.89	3,554.52	40.99	99.51	195.81
Gross income (in 1,000 USD)	82.27	116.77	43.09	66.79	103.85
Share of women	0.52	0.36	0.50	0.50	1.00
Age	37.53	13.63	27.00	33.00	45.00
First-time buyers	0.40				
Panel B: Final sales of repeat sales sequences (N = 7,160)					
# of buyers	1.45	0.52	1.00	1.00	2.00
All buyers from same HH	0.82				
Adults in HH	1.84	0.87	1.00	2.00	2.00
Any children in HH	0.24				
Any buyers with university degree	0.29				
Gross wealth (in 1,000 USD)	202.85	2,369.01	30.97	83.63	165.24
Gross income (in 1,000 USD)	72.41	70.80	38.58	58.02	92.34
Share of women	0.54	0.38	0.00	0.50	1.00
Age	35.91	13.32	26.00	31.00	43.00
First-time buyers	0.48				

Notes: The table presents summary statistics of the buyers. Panel A presents statistics for the pooled sample, and Panel B presents statistics for the repeat sales sample, reported for the buyers in the final sales of the repeat sales sequences. Variables are sorted by the order of which they appear in the buyer heterogeneity results (see Figure 2). *# of buyers* denotes the number of buyers. *All buyers from same HH* is an indicator taking the value one if all buyers belongs to the same household. *Adults in HH* denotes the minimum number of adults (age≥18) in the households to which the buyers belong. *Any children in HH* is an indicator taking the value one if there are any children in the households to which the buyers belong. *Any buyers with university degree* is an indicator taking the value one if any of the buyers hold a university degree. *Gross wealth* is the total gross wealth among the buyers, and *gross income* is the total gross income among the buyers, both adjusted to 2015 prices using the consumer price index and converted to USD at the January 2, 2023 exchange rate (NOK/USD=9.8413). *Share of women* denotes the share of women among the buyers: when zero this indicates that all buyers are men, and when one then all buyers are women. *Age* denotes the average age among the buyers. *First-time buyers* is an indicator taking the value one if none of the buyers already own a housing unit, based on the taxable value of their primary home.

Figure A.1: Distributions of the days between last showing and sale

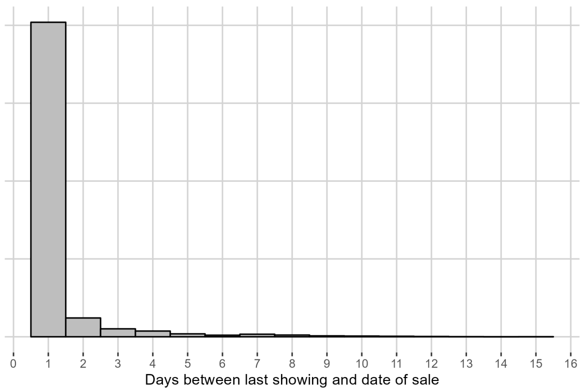

(A) Oslo

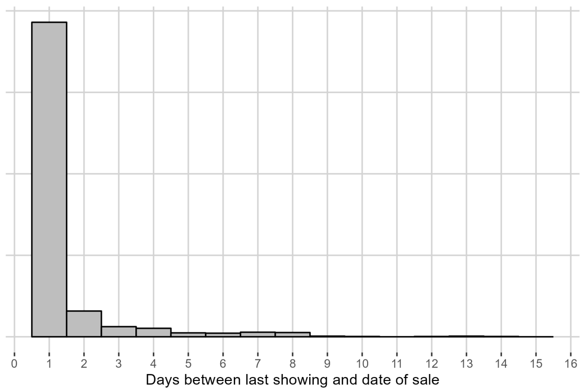

(B) Lillestrøm

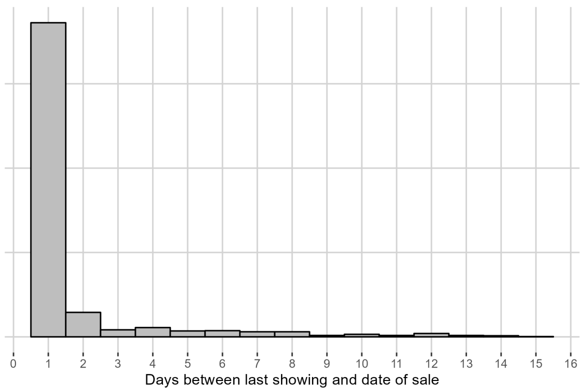

(C) Drammen

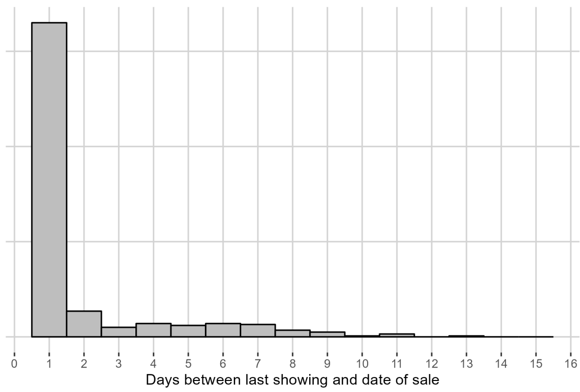

(D) Skien

Notes: The figures show distributions of the number of days between the last showing and the date of sale across four different municipalities. We chose these four municipalities to display how these distributions differ.

Figure A.2: Distributions of the days between first and last showing

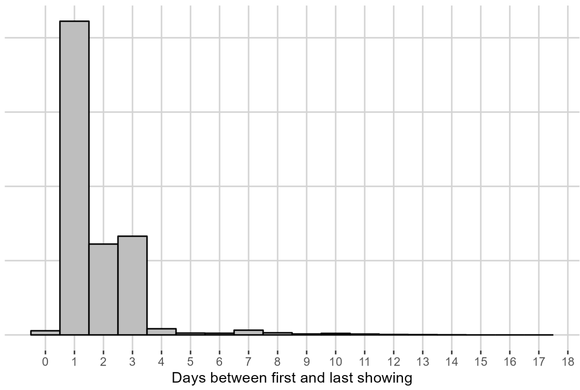

(A) Oslo

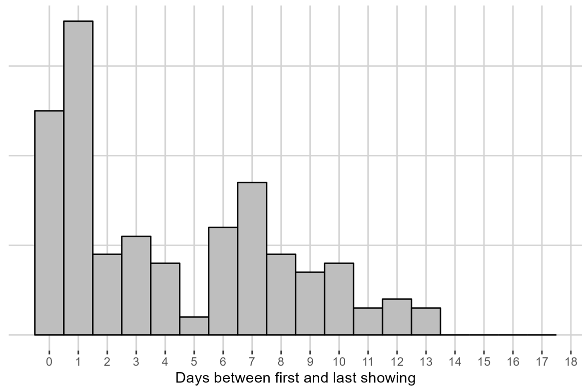

(B) Lillestrøm

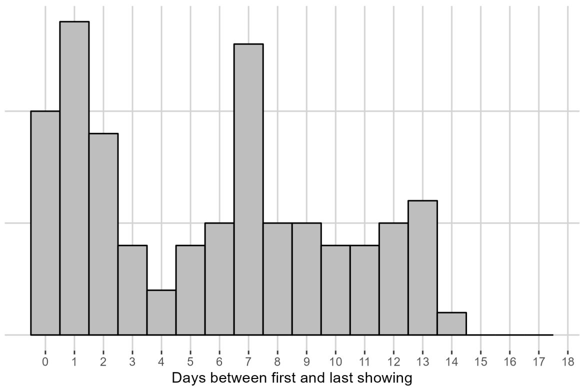

(C) Drammen

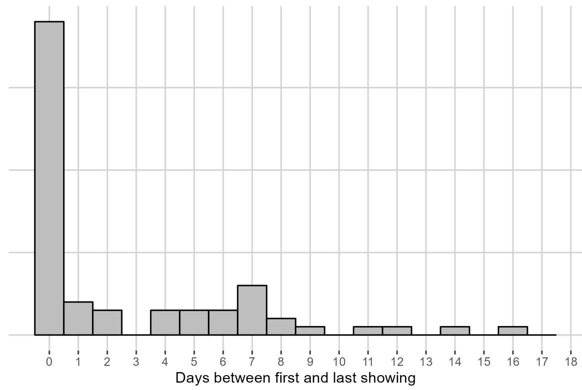

(D) Skien

Notes: The figures show distributions of the number of days between the first and the last showing for a subsample comprising sales with exactly two showings across four different municipalities. We chose these four municipalities to display how these distributions differ.

Figure A.3: Distributions at which days of the week showings are conducted

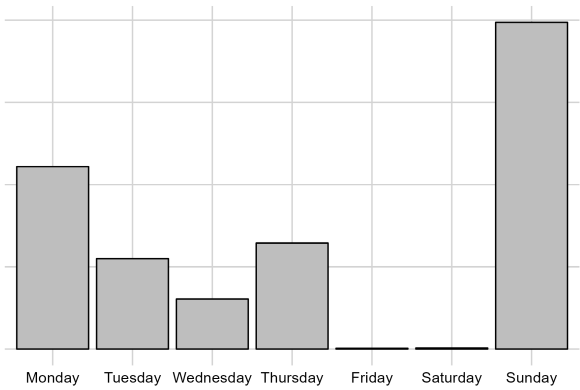

(A) Oslo

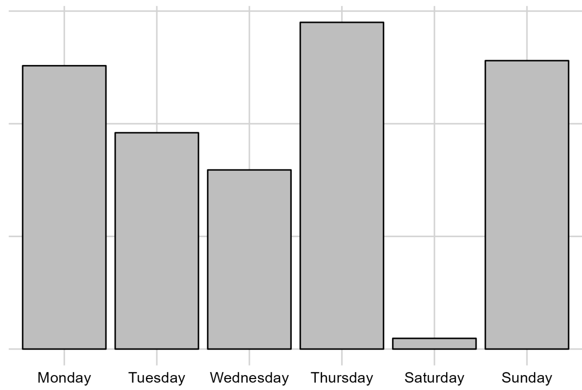

(B) Lillestrøm

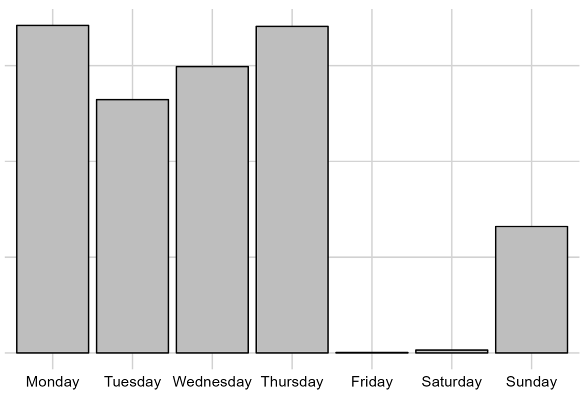

(C) Drammen

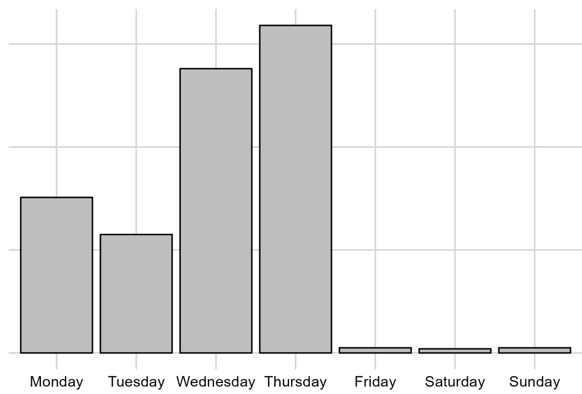

(D) Skien

Notes: The figures show which days-in-the-week showings are conducted across four different municipalities. These are based on the first and last showing dates. If there are more showings these are not included. In our supplementary sample, more than 98 percent of sales have either one or two showings. We chose these four municipalities to display how these distributions differ.

Table A.4: Weather distributions, 2015-2023, by season

	10 pct.	20 pct.	30 pct.	40 pct.	50 pct.	60 pct.	70 pct.	80 pct.	90 pct.
Spring									
Cloud cover	0.01	0.09	0.25	0.47	0.67	0.86	0.98	1.00	1.00
Precipitation	0.00	0.00	0.00	0.00	0.00	0.06	0.54	2.41	6.80
Max temp.	3.74	6.20	7.94	9.54	11.04	12.49	14.29	16.51	19.35
Wind speed	1.09	1.51	1.91	2.24	2.60	2.99	3.36	3.84	4.59
Summer									
Cloud cover	0.01	0.09	0.21	0.35	0.51	0.66	0.80	0.93	1.00
Precipitation	0.00	0.00	0.00	0.00	0.03	0.30	1.48	4.34	10.20
Max temp.	17.60	18.83	19.75	20.56	21.37	22.23	23.22	24.59	26.76
Wind speed	1.02	1.51	1.87	2.20	2.54	2.88	3.27	3.74	4.61
Fall									
Cloud cover	0.03	0.22	0.50	0.69	0.84	0.94	0.99	1.00	1.00
Precipitation	0.00	0.00	0.00	0.00	0.20	0.91	2.70	6.07	12.14
Max temp.	2.89	5.90	7.87	9.66	11.35	12.98	14.43	16.05	18.34
Wind speed	0.97	1.37	1.75	2.15	2.52	2.93	3.34	3.90	4.85

Notes: The table shows the weather distributions in our 13 municipalities in the period 2015-2023. All weather variables are extracted from gridded map data to zip codes. Cloud cover is a share between 0 (clear sky) and 1 (fully clouded), precipitation is in millimeters, max temperature is in Celsius, and wind speed is in meters per second.

Figure A.4: Observed weather distributions in south-east Norway, 2015-2023

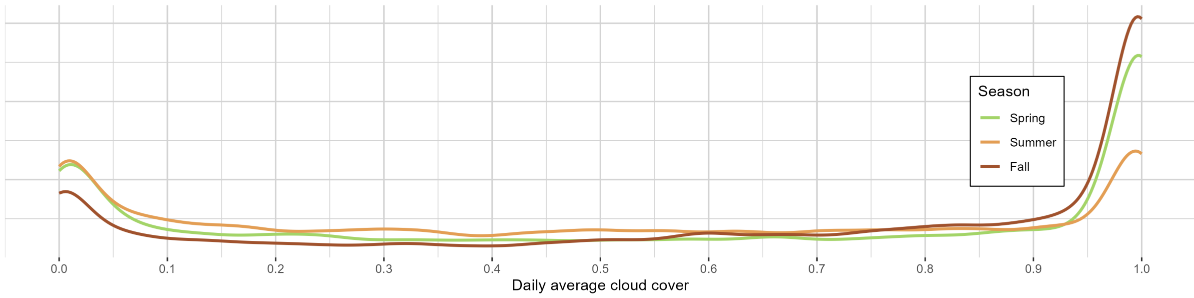

(A) Daily average cloud cover, from 0 (clear sky) to 1 (fully covered)

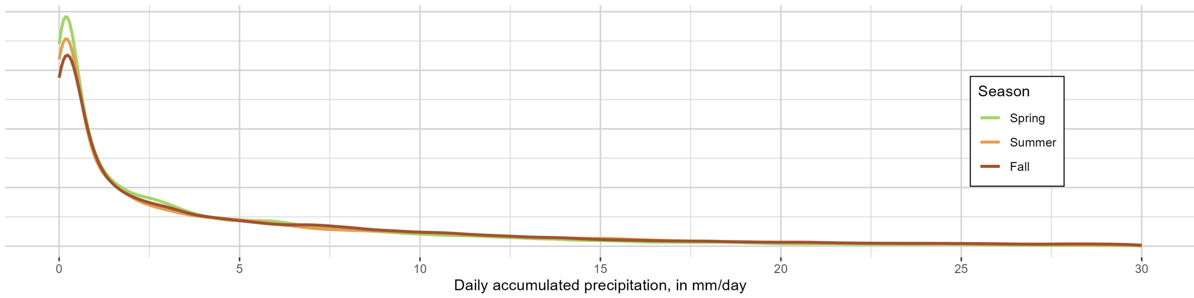

(B) Daily accumulated precipitation, in millimeters

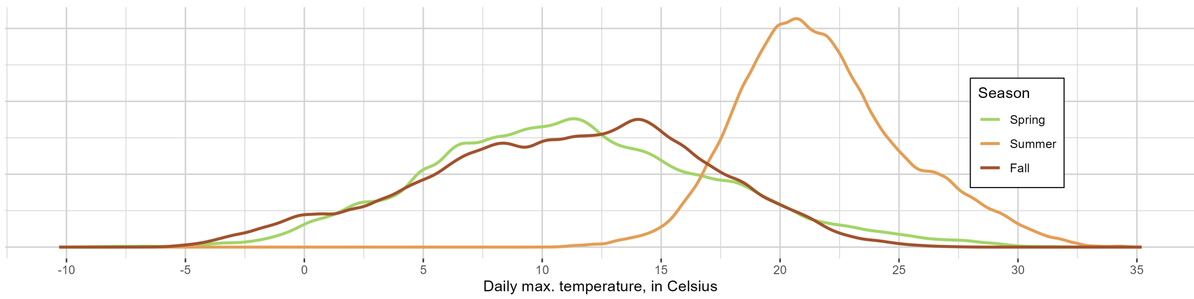

(C) Daily maximum temperature, in Celsius

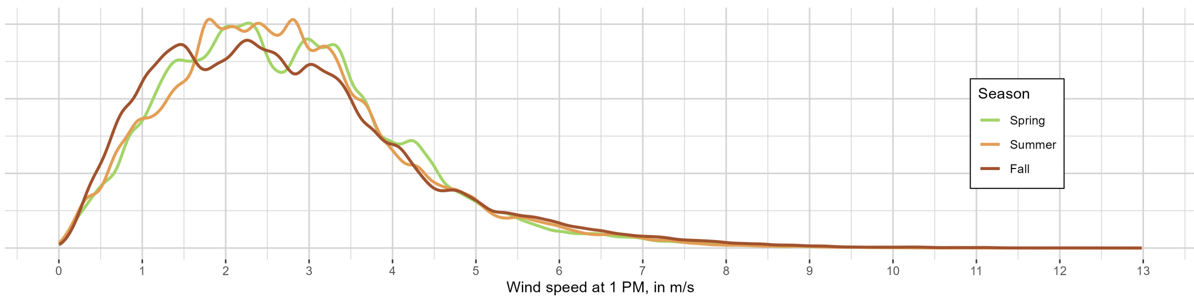

(D) Wind speed at 1:00 PM, in meters per second

Notes: The figure presents distributions of the weather variables using Gaussian kernels. Precipitation in panel (B) is zoomed in at (0,30) and does not include the boundaries.

B Additional results

B.1 Effect of cloud cover at showings on home prices - Robustness and sensitivity

B.1.1 Fixed effects and three-way interactions

We test the sensitivity of our results from changing the main specification. Results are provided in Table B.1. For the cross-sectional results, we start by keeping the cloud cover, precipitation, the seasonal controls for these, hours of night, and all fixed effects, and leaving out interactions and all other controls. The pooling effect is at the same level is before, with a MEM of -0.0126 but no longer significant. When adding the controls while keeping the interactions out, the MEM is -0.0127 and now significant. From these two estimations we are not able to distinguish between seasons because we omit the interactions. Next, we leave out the controls but include the interactions, resulting in a bit smaller MEMs that are not significant, also when calculating the MEM by seasons. It is evident that compositional factors are important for reducing the standard error. We also try including a three-way interaction between cloud cover, precipitation, and hours of night to our main specification, but the MEMs are more or less the same as before. Finally, we truncated the (log of) the hours of night to zero for all seasons except for the fall, following the approach by Kamstra et al. (2003) who argue that a similarly constructed variable should account for seasonal affective disorder (SAD), or at least some sort of seasonal depression. Using this truncated variable gives a much stronger spring MEM at -0.0202 significant at the 1% level, the summer effect is a bit lower at -0.0195 and significant at the 5% level, while fall effect is still not significant and the pooling effect remains the same as before.

We repeat the sensitivity testing for the repeat sales estimates, first by including the three-way interaction, and second by truncating the hours of night. Results from the first estimation suggest that the effects are very similar except for the fall in which the MEM is found to be -0.0347 significant at the 1% level. Thus, allowing for even more flexibility for the precipitation to be accounted for in relation to the measures of cloud cover, we find an even stronger fall effect. Results from the second estimation suggest that there is a pooling effect which is a bit lower than previously found (-0.0177, significant at the 5% level), but we now let daytime length only to be accounted for in the fall, making all other estimates of MEMs to not be significant although keeping at approximately the same magnitude. We do not prefer truncating the hours of night in this way. Cloud cover is captured in a short time period of four days, mitigating concerns about seasonal depression affecting our results. Seasonal depression should not affect prices in relation with cloud cover in such a short time interval. We rather believe that exposure to cloud cover in the short time interval should matter differently in the spring, summer and fall, so that only considering hours of night during the fall imposes an assumption of seasonality that is too strict.

B.1.2 Adjusting the moving average window

Our models use the average weather of the four days prior to the sale. The moving average window is empirically and institutionally motivated, and we test the sensitivity to this window by both reducing it to two days and increasing it to six day. We abstain from increasing the moving average window further because it will start resembling the seasonal controls. Results are provided in Table B.2. For the cross-sectional results, we find that reducing the window increases the pooling effect and reduces the summer effect, but all MEMs are significant at the 1% level except for the fall which is still not significant. Increasing the window gives approximately the same pooling result as for the reduced window, significant at the 10% level, but a larger increase in the summer effect (0.0350, significant at the 1% level). For the repeat sales results, when reducing the window the MEMs are smaller and no longer significant, and for the increased window the MEMs are approximately the same as before. However, increasing the window and imposing the holding time restriction as before gives only a significant fall effect at the 10% level (-0.280).

When changing the moving average window, we allow for changes in two dimensions. First, when increasing (reducing) the window, we allow the moving average to be more (less) correlated with the seasonal controls, which allows for biases in unknown directions. Second, taking the average over more days adds noise to the weather that is relevant to include, while reducing the window makes us not measure the weather at important days. Meaning, by increasing the window we may add noise and irrelevant information while also amplifying the estimates, resulting in stronger but less significant estimates. And by reducing the window we reduce both the possibility of measuring important weather and correlation with seasonal controls, thus we no longer estimate the effect we seek to investigate.

B.1.3 Shorter TOM

As an alternative to changing the moving average window, we can rather increase the likelihood of us measuring the cloud cover at the relevant showing days by imposing constraints on the time-on-market (TOM). When reducing the TOM in the sample, we ensure that public showings are conducted closer in time to the sale, thus it becomes more likely that the showings are sometime during the four days prior to the sale. The maximum TOM in the sample is 22 days and the minimum is 7 days, and we choose to restrict at 18, 14, and 10 days between listing and sale. The results, provided in Table B.3, indicate that the shorter TOM, the stronger the effect becomes. This is consistent with the fact that we are more likely to measure the cloud cover at the relevant days. The most notable result is found for the repeat sales estimates when restricting the sample to consist of transactions with at most 10 days TOM: the pooling MEM is -0.0345, significant at the 5% level. This estimate is almost twice the size of the estimates found in Table 2.

Table B.1: Sensitivity to alternative controls, fixed effects, and three-way interactions

	Cross-section					Unit FE	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Marginal effects at the means, pooled and by season							
Pooled	-0.0126	-0.0127**	-0.0082	-0.0122*	-0.0123*	-0.0202**	-0.0177**
	(0.0118)	(0.0058)	(0.0131)	(0.0064)	(0.0065)	(0.0090)	(0.0090)
Spring			-0.0113	-0.0138**	-0.0202***	-0.0200**	-0.0170
			(0.0125)	(0.0061)	(0.0074)	(0.0086)	(0.0104)
Summer			-0.0208	-0.0245***	-0.0195**	-0.0072	-0.0148
			(0.0171)	(0.0088)	(0.0079)	(0.0121)	(0.0110)
Fall			0.0082	0.0013	0.0048	-0.0347***	-0.0215
			(0.0185)	(0.0095)	(0.0102)	(0.0131)	(0.0145)
N	55,818	52,906	55,818	52,906	52,906	12,170	12,170
Adj. R. sq.	0.4102	0.8728	0.4102	0.8728	0.8728	0.9593	0.9594
Controls		✓		✓	✓	✓	✓
Interactions			Two-way	Three-way	Two-way	Three-way	Two-way
SAD, not HoN					✓		✓
Fixed effects	MYQ	MYQ	MYQ	MYQ	MYQ	YQ	YQ
Zip code FE	✓	✓	✓	✓	✓		

Notes: The table shows marginal effects at the means from regressing log of the selling price on the cloud cover, in which the selling price includes common debt if any. Weather variables are the moving averages of the four days prior to the sale. All variables are in natural logarithms. HoN denotes the hours of night. MYQ and YQ denotes the municipality-by-year-by-quarter and year-by-quarter fixed effects, respectively. The weather controls maximum temperature and wind speed and their seasonal monthly averages. Other controls are age, mortgage rate, size (square meters), an indicator taking the value of one if the housing unit is an apartment, the interaction between size and the apartment indicator, an indicator taking the value of one if the housing unit is an owner-occupier unit, and an interaction between lot size (square meters) and a non-apartment indicator. The three-way interaction is an interaction between cloud cover, precipitation, and hours of night. SAD is an alternative to the SAD variable presented by Kamstra et al. (2003): our SAD is the natural logarithm of the hours of night truncated at 0 for the spring and summer. Heteroskedasticity-robust standard errors are given in parenthesis. Significance: * $p<0.1$, ** $p<0.05$, *** $p<0.01$.

Table B.2: Sensitivity to changing the moving average window

	MA(2)			MA(6)		
	(1)	(2)	(3)	(4)	(5)	(6)
Marginal effects at the means, pooled and by season						
Pooled	-0.0138***	-0.0099	-0.0105	-0.0136*	-0.0191*	-0.0157
	(0.0052)	(0.0070)	(0.0070)	(0.0072)	(0.0104)	(0.0102)
Spring	-0.0136***	-0.0095	-0.0092	-0.0162**	-0.0197*	-0.0161
	(0.0049)	(0.0066)	(0.0066)	(0.0069)	(0.0101)	(0.0099)
Summer	-0.0204***	-0.0078	-0.0072	-0.0350***	-0.0034	-0.0028
	(0.0067)	(0.0091)	(0.0088)	(0.0096)	(0.0133)	(0.0130)
Fall	-0.0073	-0.0125	-0.0155	0.0109	-0.0341**	-0.0280*
	(0.0072)	(0.0102)	(0.0103)	(0.0103)	(0.0148)	(0.0147)
N	52,906	12,170	11,412	52,906	12,170	11,412
Adj. R. sq.	0.8728	0.9593	0.9621	0.8728	0.9593	0.9621
Fixed effects	MYQ	YQ	YQ	MYQ	YQ	YQ
Zip code FE	✓			✓		
Unit FE		✓	✓		✓	✓
Holding time (weeks)			≥52			≥52

Notes: The table shows marginal effects at the means from regressing log of the selling price on the cloud cover, in which the selling price includes common debt if any. Weather variables in the first three columns are the moving averages of the two days prior to the sale (MA(2)), and weather variables in the last three columns are the moving averages of the six days prior to the sale (MA(6)). All variables are in natural logarithms. MYQ and YQ denotes the municipality-by-year-by-quarter and year-by-quarter fixed effects, respectively. All regressions include the weather controls maximum temperature and wind speed and their seasonal monthly averages. All regressions control for age and mortgage rate. Cross-section regressions control for attributes: size (square meters), an indicator taking the value of one if the housing unit is an apartment, the interaction between size and the apartment indicator, an indicator taking the value of one if the housing unit is an owner-occupier unit, and an interaction between lot size (square meters) and a non-apartment indicator. Heteroskedasticity-robust standard errors are given in parenthesis. Significance: * p<0.1, ** p<0.05, *** p<0.01.

Table B.3: Sensitivity to restricting TOM in the sample

	Cross-section			Unit FE		
TOM	≤18	≤14	≤10	≤18	≤14	≤10
Marginal effects at the means, pooled and by season						
Pooled	-0.0152**	-0.0167**	-0.0179**	-0.0193**	-0.0214**	-0.0345**
	(0.0066)	(0.0068)	(0.0081)	(0.0094)	(0.0098)	(0.0135)
Spring	-0.0175***	-0.0182***	-0.0203***	-0.0197**	-0.0217**	-0.0327**
	(0.0063)	(0.0065)	(0.0077)	(0.0091)	(0.0095)	(0.0129)
Summer	-0.0330***	-0.0336***	-0.0345***	-0.0129	-0.0140	-0.0205
	(0.0087)	(0.0090)	(0.0105)	(0.0123)	(0.0130)	(0.0176)
Fall	0.0053	0.0022	0.0017	-0.0251*	-0.0285**	-0.0506**
	(0.0093)	(0.0097)	(0.0114)	(0.0133)	(0.0141)	(0.0198)
N	50,359	46,791	33,032	11,104	9,754	5,366
Adj. R. sq.	0.8733	0.8739	0.8736	0.9586	0.9589	0.9556
Fixed effects	MYQ	MYQ	MYQ	YQ	YQ	YQ
Zip code FE	✓	✓	✓			

Notes: The table shows marginal effects at the means from regressing log of the selling price on the cloud cover when restricting the time-on-market, in which the selling price includes common debt if any. Weather variables are the moving averages of the four days prior to the sale. All variables are in natural logarithms. MYQ and YQ denote the municipality-by-year-by-quarter and year-by-quarter fixed effects, respectively. All regressions include the weather controls maximum temperature and wind speed and their seasonal monthly averages. All regressions control for age and mortgage rate. Cross-section regressions control for attributes: size (square meters), an indicator taking the value of one if the housing unit is an apartment, the interaction between size and the apartment indicator, an indicator taking the value of one if the housing unit is an owner-occupier unit, and an interaction between lot size (square meters) and a non-apartment indicator. Heteroskedasticity-robust standard errors are given in parenthesis. Significance: * p<0.1, ** p<0.05, *** p<0.01.

B.2 Effect of cloud cover at showings and day of sale on home prices

Table B.4: Marginal effects of cloud cover at showings and day of sale on log of home prices (evaluated at means; MEMs), unit FE, restricted

	Non-residualized		Residualized	
	(1)	(2)	(3)	(4)
Panel A: Day of the sale				
Pooled	-0.0018	-0.0008	0.0011	-0.0015
	(0.0061)	(0.0062)	(0.0053)	(0.0054)
Spring	-0.0019	-0.0012	0.0002	-0.0022
	(0.0054)	(0.0055)	(0.0050)	(0.0051)
Summer	-0.0033	-0.0040	-0.0018	-0.0047
	(0.0074)	(0.0076)	(0.0069)	(0.0070)
Fall	-0.0003	0.0031	0.0052	0.0026
	(0.0091)	(0.0093)	(0.0078)	(0.0079)
Panel B: Days prior to the sale, MA(4)				
Pooled		-0.0162*		-0.0165*
		(0.0090)		(0.0090)
Spring		-0.0157*		-0.0161*
		(0.0087)		(0.0087)
Summer		-0.0101		-0.0114
		(0.0116)		(0.0116)
Fall		-0.0229*		-0.0222*
		(0.0128)		(0.0128)
N	10,998	10,998	10,998	10,998
Adj. R. sq.	0.9623	0.9623	0.9623	0.9623
F stat.	3.2959	2.7461	3.4426	2.8241
Fixed effects	YQ	YQ	YQ	YQ
Holding time (weeks)	≥52	≥52	≥52	≥52

Notes: The table shows marginal effects at the means from regressing log of the selling price on the cloud cover, in which the selling price includes common debt if any. All results are from estimating specification (1) with unit fixed effects, and restricting the sample to holding times of between 52 and 520 weeks. Column (1) estimates the model with the weather at the day of the sale, while column (2) adds the moving averages as well. Columns (3) and (4) substitutes the raw weather at the day of the sale with residualized weather. YQ denotes the year-by-quarter fixed effects. All regressions include the weather controls maximum temperature and wind speed and their seasonal monthly averages. Heteroskedasticity-robust standard errors are given in parenthesis. Significance: * p<0.1, ** p<0.05, *** p<0.01.

B.3 Heterogeneity of cloud cover at showings

Table B.5: Segmented regressions on apartments and non-apartments

	Apartments			Non-apartments		
	(1)	(2)	(3)	(4)	(5)	(6)
Marginal effects at the means, pooled and by season						
Pooled	-0.0053	-0.0229**	-0.0189**	-0.0132	-0.0139	-0.0040
	(0.0069)	(0.0091)	(0.0090)	(0.0137)	(0.0291)	(0.0284)
Spring	-0.0071	-0.0244***	-0.0204**	-0.0141	-0.0070	0.0061
	(0.0066)	(0.0087)	(0.0086)	(0.0132)	(0.0284)	(0.0280)
Summer	-0.0155*	-0.0233**	-0.0214*	-0.0272	0.0196	0.0365
	(0.0090)	(0.0116)	(0.0114)	(0.0185)	(0.0402)	(0.0386)
Fall	0.0070	-0.0205	-0.0146	0.0019	-0.0558	-0.0570
	(0.0096)	(0.0129)	(0.0128)	(0.0193)	(0.0396)	(0.0389)
N	38,991	10,172	9,561	13,908	1,996	1,849
Adj. R. sq.	0.8728	0.9631	0.9659	0.8770	0.9461	0.9491
Fixed effects	MYQ	YQ	YQ	MYQ	YQ	YQ
Zip code FE	✓			✓		
Unit FE		✓	✓		✓	✓
Holding time (weeks)			≥52			≥52

Notes: The table shows results of regressing log of the selling price on the cloud cover, in which the selling price includes common debt if any. The results are from subsamples of apartments and non-apartments. The effects are evaluated at the sample means of the full sample, as well as by season, without conditioning on segmentation. Weather variables are the moving averages of the four days prior to the sale. All variables are in natural logarithms. Fixed effects are abbreviated as follows: M is municipality, Y is year, Q is quarter, and DW is day-of-the-week. Thus, MYQ denotes the municipality-by-year-by-quarter fixed effects. All regressions include the weather controls maximum temperature and wind speed and their seasonal monthly averages. All regressions control for age and mortgage rate. For cross-section regressions covariates are different between the apartments and non-apartments subsamples. The apartments regression includes an indicator taking the value of one if the housing unit is an owner-occupier unit, and the non-apartments regressions include the lot size (square meters) and housing unit category indicators; row house, semi-detached, and with detached being the omitted category. Heteroskedasticity-robust standard errors are given in parenthesis. Significance: * p<0.1, ** p<0.05, *** p<0.01.

Table B.6: Buyer heterogeneity - buyers and households

	Buyers		Adults in HH		Children in HH	
	1	>1	1	>1	0	>0
Marginal effects at the means, pooled and by season						
Pooled	-0.0109	-0.0191**	-0.0191	-0.0108	-0.0126	-0.0197
	(0.0085)	(0.0094)	(0.0130)	(0.0087)	(0.0086)	(0.0130)
Spring	-0.0119	-0.0206**	-0.0198	-0.0123	-0.0138*	-0.0223*
	(0.0081)	(0.0090)	(0.0125)	(0.0083)	(0.0082)	(0.0125)
Summer	-0.0183*	-0.0404***	-0.0199	-0.0317***	-0.0211*	-0.0489***
	(0.0111)	(0.0125)	(0.0170)	(0.0115)	(0.0113)	(0.0173)
Fall	-0.0021	0.0040	-0.0174	0.0121	-0.0026	0.0127
	(0.0121)	(0.0131)	(0.0184)	(0.0121)	(0.0121)	(0.0181)
N	25,917	25,926	11,898	30,945	28,798	14,045
Adj. R. sq.	0.8550	0.8648	0.8605	0.8765	0.8642	0.8896
Fixed effects	MYQ	MYQ	MYQ	MYQ	MYQ	MYQ
Zip code FE	✓	✓	✓	✓	✓	✓

Notes: The table shows marginal effects at the means from regressing log of the selling price on the cloud cover, in which the selling price includes common debt if any. The samples are segmented by the characteristics of the buyers. *Buyers* denotes segmenting by the number of buyers. *Adults in HH* denotes segmenting by the number of adults in the household, focusing on transactions in which all buyers belong to the same household. *Children in HH* denotes segmenting by the number of children in the household, defined as those up to 17 years old, focusing on transactions in which all buyers belong to the same household. The effects are evaluated at the sample means of the full sample, as well as by season, without conditioning on segmentation. Weather variables are the moving averages of the four days prior to the sale. All variables are in natural logarithms. MYQ denotes the municipality-by-year-by-quarter fixed effects. All regressions include the weather controls maximum temperature and wind speed and their seasonal monthly averages. All regressions control for age and mortgage rate, size (square meters), an indicator taking the value of one if the housing unit is an apartment, the interaction between size and the apartment indicator, an indicator taking the value of one if the housing unit is an owner-occupier unit, and an interaction between lot size (square meters) and a non-apartment indicator. Heteroskedasticity-robust standard errors are given in parenthesis. Significance: * p<0.1, ** p<0.05, *** p<0.01.

Table B.7: Buyer heterogeneity - education, wealth and income

	University degree		Wealth		Income	
	No	Yes	$\leq$ median	>median	$\leq$ median	>median
Marginal effects at the means, pooled and by season						
Pooled	-0.0117	-0.0143*	-0.0058	-0.0142	-0.0192**	-0.0117
	(0.0116)	(0.0077)	(0.0075)	(0.0096)	(0.0083)	(0.0095)
Spring	-0.0123	-0.0161**	-0.0064	-0.0161*	-0.0182**	-0.0145
	(0.0111)	(0.0074)	(0.0072)	(0.0092)	(0.0079)	(0.0091)
Summer	-0.0288*	-0.0302***	-0.0086	-0.0402***	-0.0235**	-0.0367***
	(0.0156)	(0.0101)	(0.0098)	(0.0128)	(0.0107)	(0.0127)
Fall	0.0062	0.0036	-0.0024	0.0140	-0.0160	0.0167
	(0.0162)	(0.0108)	(0.0107)	(0.0135)	(0.0118)	(0.0133)
N	16,188	34,632	25,682	25,623	25,872	25,435
Adj. R. sq.	0.8352	0.8783	0.8292	0.8600	0.8371	0.8609
Fixed effects	MYQ	MYQ	MYQ	MYQ	MYQ	MYQ
Zip code FE	✓	✓	✓	✓	✓	✓

Notes: The table shows marginal effects at the means from regressing log of the selling price on the cloud cover, in which the selling price includes common debt if any. The samples are segmented by the characteristics of the buyers. *University degree* denotes segmenting by whether any of the buyer have a university degree. *Wealth* denotes segmenting by the total gross wealth among the buyers. *Income* denotes segmenting by the total gross income among the buyers. Wealth and income are deflated to 2015 price levels using the consumer price index. The effects are evaluated at the sample means of the full sample, as well as by season, without conditioning on segmentation. Weather variables are the moving averages of the four days prior to the sale. All variables are in natural logarithms. MYQ denotes the municipality-by-year-by-quarter fixed effects. All regressions include the weather controls maximum temperature and wind speed and their seasonal monthly averages. All regressions control for age and mortgage rate, size (square meters), an indicator taking the value of one if the housing unit is an apartment, the interaction between size and the apartment indicator, an indicator taking the value of one if the housing unit is an owner-occupier unit, and an interaction between lot size (square meters) and a non-apartment indicator. Heteroskedasticity-robust standard errors are given in parenthesis. Significance: * $p<0.1$, ** $p<0.05$, *** $p<0.01$.

Table B.8: Buyer heterogeneity - gender, age and first-time buyers

	Gender (1 buyer)		Age		First-time buyers	
	Male	Female	$\leq$ median	>median	No	Yes
Marginal effects at the means, pooled and by season						
Pooled	-0.0070	-0.0120	-0.0123	-0.0157	-0.0170*	-0.0068
	(0.0127)	(0.0117)	(0.0080)	(0.0101)	(0.0088)	(0.0085)
Spring	-0.0085	-0.0130	-0.0130*	-0.0173*	-0.0189**	-0.0074
	(0.0121)	(0.0111)	(0.0076)	(0.0096)	(0.0084)	(0.0081)
Summer	-0.0199	-0.0150	-0.0208**	-0.0377***	-0.0429***	-0.0078
	(0.0167)	(0.0150)	(0.0104)	(0.0134)	(0.0116)	(0.0112)
Fall	0.0079	-0.0076	-0.0028	0.0082	0.0112	-0.0048
	(0.0181)	(0.0164)	(0.0114)	(0.0139)	(0.0123)	(0.0119)
N	12,103	13,423	25,998	25,308	30,635	20,667
Adj. R. sq.	0.8555	0.8571	0.8803	0.8724	0.8679	0.8519
Fixed effects	MYQ	MYQ	MYQ	MYQ	MYQ	MYQ
Zip code FE	✓	✓	✓	✓	✓	✓

Notes: The table shows marginal effects at the means from regressing log of the selling price on the cloud cover, in which the selling price includes common debt if any. The samples are segmented by the characteristics of the buyers. *Gender* denotes segmenting by the gender of buyers, focusing on transactions with one buyer. *Age* denotes segmenting by the average age among the buyers. *First-time buyers* denotes segmenting by whether none of the buyers already own a housing unit, based on the taxable value of their primary home. The effects are evaluated at the sample means of the full sample, as well as by season, without conditioning on segmentation. Weather variables are the moving averages of the four days prior to the sale. All variables are in natural logarithms. MYQ denotes the municipality-by-year-by-quarter fixed effects. All regressions include the weather controls maximum temperature and wind speed and their seasonal monthly averages. All regressions control for age and mortgage rate, size (square meters), an indicator taking the value of one if the housing unit is an apartment, the interaction between size and the apartment indicator, an indicator taking the value of one if the housing unit is an owner-occupier unit, and an interaction between lot size (square meters) and a non-apartment indicator. Heteroskedasticity-robust standard errors are given in parenthesis. Significance: * p<0.1, ** p<0.05, *** p<0.01.

Table B.9: Buyer heterogeneity, repeat sales - buyers and households

	Buyers		Adults in HH		Children in HH	
	1	>1	1	>1	0	>0
Marginal effects at the means, pooled and by season						
Pooled	-0.0328***	-0.0029	-0.0360**	-0.0180	-0.0221*	-0.0286
	(0.0114)	(0.0146)	(0.0175)	(0.0124)	(0.0117)	(0.0213)
Spring	-0.0307***	-0.0068	-0.0328**	-0.0190	-0.0227**	-0.0280
	(0.0111)	(0.0139)	(0.0167)	(0.0120)	(0.0112)	(0.0205)
Summer	-0.0225	-0.0012	-0.0170	-0.0108	-0.0160	-0.0112
	(0.0148)	(0.0189)	(0.0232)	(0.0160)	(0.0150)	(0.0277)
Fall	-0.0457***	0.0002	-0.0587**	-0.0240	-0.0274*	-0.0466
	(0.0161)	(0.0204)	(0.0244)	(0.0177)	(0.0166)	(0.0300)
N	6,709	5,311	3,091	6,557	7,163	2,487
Adj. R. sq.	0.9574	0.9511	0.9565	0.9593	0.9574	0.9612
Fixed effects	YQ	YQ	YQ	YQ	YQ	YQ
Unit FE	✓	✓	✓	✓	✓	✓

Notes: The table shows marginal effects at the means from regressing log of the selling price on the cloud cover, in which the selling price includes common debt if any. The samples are segmented by the characteristics of the buyers in the last sale within the repeat sales sequence. *Buyers* denotes segmenting by the number of buyers. *Adults in HH* denotes segmenting by the number of adults in the household, focusing on transactions in which all buyers belong to the same household. *Children in HH* denotes segmenting by the number of children in the household, defined as those up to 17 years old, focusing on transactions in which all buyers belong to the same household. The effects are evaluated at the sample means of the full sample, as well as by season, without conditioning on segmentation. Weather variables are the moving averages of the four days prior to the sale. All variables are in natural logarithms. YQ denotes the year-by-quarter fixed effects. All regressions include the weather controls maximum temperature and wind speed and their seasonal monthly averages. All regressions control for age and mortgage rate.
Heteroskedasticity-robust standard errors are given in parenthesis. Significance: * p<0.1, ** p<0.05, *** p<0.01.

Table B.10: Buyer heterogeneity, repeat sales - education, wealth and income

	University degree		Wealth		Income	
	No	Yes	$\leq$ median	>median	$\leq$ median	>median
Marginal effects at the means, pooled and by season						
Pooled	-0.0148	-0.0175	-0.0325***	0.0007	-0.0176	-0.0201
	(0.0169)	(0.0109)	(0.0114)	(0.0149)	(0.0113)	(0.0154)
Spring	-0.0124	-0.0194*	-0.0325***	-0.0010	-0.0166	-0.0228
	(0.0167)	(0.0104)	(0.0110)	(0.0143)	(0.0109)	(0.0147)
Summer	-0.0288	-0.0036	-0.0407***	0.0252	-0.0131	-0.0081
	(0.0227)	(0.0140)	(0.0145)	(0.0198)	(0.0143)	(0.0201)
Fall	-0.0037	-0.0291*	-0.0243	-0.0219	-0.0235	-0.0286
	(0.0233)	(0.0153)	(0.0161)	(0.0206)	(0.0158)	(0.0216)
N	3,376	8,402	6,801	5,095	6,837	5,059
Adj. R. sq.	0.9516	0.9578	0.9360	0.9564	0.9517	0.9502
Fixed effects	YQ	YQ	YQ	YQ	YQ	YQ
Unit FE	✓	✓	✓	✓	✓	✓

Notes: The table shows marginal effects at the means from regressing log of the selling price on the cloud cover, in which the selling price includes common debt if any. The samples are segmented by the characteristics of the buyers in the last sale within the repeat sales sequence. *University degree* denotes segmenting by whether any of the buyer have a university degree. *Wealth* denotes segmenting by the total gross wealth among the buyers. *Income* denotes segmenting by the total gross income among the buyers. Wealth and income are deflated to 2015 price levels using the consumer price index. The effects are evaluated at the sample means of the full sample, as well as by season, without conditioning on segmentation. Weather variables are the moving averages of the four days prior to the sale. All variables are in natural logarithms. YQ denotes the year-by-quarter fixed effects. All regressions include the weather controls maximum temperature and wind speed and their seasonal monthly averages. All regressions control for age and mortgage rate. Heteroskedasticity-robust standard errors are given in parenthesis. Significance: * $p<0.1$, ** $p<0.05$, *** $p<0.01$.

Table B.11: Buyer heterogeneity, repeat sales - gender, age and first-time buyers

	Gender (1 buyer)		Age		First-time buyers	
	Male	Female	$\leq$ median	>median	No	Yes
Marginal effects at the means, pooled and by season						
Pooled	-0.0057	-0.0518***	-0.0164	-0.0213	-0.0112	-0.0244**
	(0.0164)	(0.0158)	(0.0117)	(0.0145)	(0.0135)	(0.0123)
Spring	-0.0065	-0.0479***	-0.0184	-0.0204	-0.0130	-0.0243**
	(0.0159)	(0.0153)	(0.0112)	(0.0141)	(0.0131)	(0.0118)
Summer	-0.0131	-0.0287	-0.0185	-0.0025	0.0046	-0.0292*
	(0.0220)	(0.0201)	(0.0153)	(0.0188)	(0.0173)	(0.0160)
Fall	0.0027	-0.0799***	-0.0117	-0.0413**	-0.0249	-0.0199
	(0.0242)	(0.0215)	(0.0168)	(0.0194)	(0.0187)	(0.0173)
N	2,856	3,769	6,822	5,074	6,250	5,646
Adj. R. sq.	0.9553	0.9590	0.9542	0.9635	0.9558	0.9504
Fixed effects	YQ	YQ	YQ	YQ	YQ	YQ
Unit FE	✓	✓	✓	✓	✓	✓

Notes: The table shows marginal effects at the means from regressing log of the selling price on the cloud cover, in which the selling price includes common debt if any. The samples are segmented by the characteristics of the buyers in the last sale within the repeat sales sequence. *Gender* denotes segmenting by the gender of buyers, focusing on transactions with one buyer. *Age* denotes segmenting by the average age among the buyers. *First-time buyers* denotes segmenting by whether none of the buyers already own a housing unit, based on the taxable value of their primary home. The effects are evaluated at the sample means of the full sample, as well as by season, without conditioning on segmentation. Weather variables are the moving averages of the four days prior to the sale. All variables are in natural logarithms. YQ denotes the year-by-quarter fixed effects. All regressions include the weather controls maximum temperature and wind speed and their seasonal monthly averages. All regressions control for age and mortgage rate. Heteroskedasticity-robust standard errors are given in parenthesis. Significance: * $p<0.1$, ** $p<0.05$, *** $p<0.01$.

B.4 Validity of cloud cover at showings, cross-sectional sample

Table B.12: Validity

Estimates			At means			
CC	CC x Prec.	CC x HoN	Pooled	Spring	Summer	Fall
Replace all						
0.02005	-0.00090	-0.00803	0.00118	0.00160	0.00375	-0.00193
(0.02353)	(0.00584)	(0.01037)	(0.00561)	(0.00517)	(0.00682)	(0.00711)
Replace cloud cover						
0.00028	-0.00006	-0.00013	-0.00007	-0.00005	-0.00003	-0.00012
(0.02754)	(0.00520)	(0.01192)	(0.00440)	(0.00444)	(0.00589)	(0.00634)
Replace precipitation						
-0.09034	0.00007	0.03733	-0.00649	-0.00790	-0.01876	0.00752
(0.00478)	(0.00506)	(0.00013)	(0.00007)	(0.00067)	(0.00041)	(0.00059)

Notes: The table shows results from re-estimating model (1) by replacing the weather realizations by random draws without replacement from the time series of moving averages. Reported coefficients are the means across repetitions, and standard deviations of the coefficient estimates are given in parenthesis. When replacing all variables, this implies replacing all weather variables including max temperature, wind speed, and all seasonal control by full realizations of those (row-wise replacement). When replacing a single weather variable, this implies replacing just that moving average. The procedure is repeated 1,000 times, and the reported coefficient estimates are the averages across the trials. Standard deviations of coefficient estimates are provided in parenthesis. All variables are in natural logarithms. Fixed effects are the municipality-by-year-by-quarter fixed effects. All regressions include the weather controls maximum temperature and wind speed and their seasonal monthly averages. All regressions control for age and mortgage rate, size (square meters), an indicator taking the value of one if the housing unit is an apartment, the interaction between size and the apartment indicator, an indicator taking the value of one if the housing unit is an owner-occupier unit, and an interaction between lot size (square meters) and a non-apartment indicator.

Table B.13: Validity, repeat sales sample

Estimates			At means			
CC	CC x Prec.	CC x HoN	All	Spring	Summer	Fall
Replace all						
0.02753	-0.00098	-0.01188	0.00002	0.00060	0.00386	-0.00454
(0.03771)	(0.00847)	(0.01631)	(0.00816)	(0.00780)	(0.01012)	(0.01034)
Replace cloud cover						
0.00021	-0.00021	0.00004	0.00011	0.00014	0.00008	0.00010
(0.04154)	(0.00745)	(0.01781)	(0.00636)	(0.00647)	(0.00854)	(0.00941)
Replace precipitation						
0.03528	-0.00022	-0.02081	-0.01085	-0.00997	-0.00345	-0.01933
(0.00682)	(0.00714)	(0.00047)	(0.00019)	(0.00095)	(0.00061)	(0.00087)

Notes: The table shows results from re-estimating model (1) by replacing the weather realizations by random draws without replacement from the time series of moving averages. Reported coefficients are the means across repetitions, and standard deviations of the coefficient estimates are given in parenthesis. When replacing all variables, this implies replacing all weather variables including max temperature, wind speed, and all seasonal control by full realizations of those (row-wise replacement). When replacing a single weather variable, this implies replacing just that moving average. The procedure is repeated 1,000 times, and the reported coefficient estimates are the averages across the trials. Standard deviations of coefficient estimates are provided in parenthesis. All variables are in natural logarithms. Fixed effects are unit and year-by-quarter fixed effects. All regressions include the weather controls maximum temperature and wind speed and their seasonal monthly averages. All regressions control for age and mortgage rate.

B.5 Removing observations with dissimilar weather at showings and day of the sale

Table B.14: Remove dissimilar cloud cover

$\lvert CC^{DOS} - CC^{MA} \rvert$	≤0.8	≤0.7	≤0.6	≤0.5	≤0.4
Marginal effects at the means, pooled and by season					
Pooled	-0.0170*	-0.0199**	-0.0267**	-0.0240**	-0.0256*
	(0.0095)	(0.0101)	(0.0112)	(0.0121)	(0.0148)
Spring	-0.0173*	-0.0208**	-0.0275**	-0.0251**	-0.0276*
	(0.0091)	(0.0098)	(0.0108)	(0.0118)	(0.0141)
Summer	-0.0150	-0.0214	-0.0218	-0.0184	-0.0120
	(0.0122)	(0.0132)	(0.0145)	(0.0156)	(0.0185)
Fall	-0.0184	-0.0171	-0.0304*	-0.0280*	-0.0369*
	(0.0133)	(0.0143)	(0.0156)	(0.0170)	(0.0213)
N	11,373	10,497	9,027	7,506	5,249
Adj. R. sq.	0.9593	0.9588	0.9582	0.9575	0.9569
Fixed effects	YQ	YQ	YQ	YQ	YQ
Unit FE	✓	✓	✓	✓	✓

Notes: The table shows marginal effects at the means from regressing log of the selling price on the cloud cover, in which the selling price includes common debt if any. Weather variables are the moving averages of the four days prior to the sale. All variables are in natural logarithms. YQ denotes the year-by-quarter fixed effects. All regressions include the weather controls maximum temperature and wind speed and their seasonal monthly averages. All regressions control for age and mortgage rate. Heteroskedasticity-robust standard errors are given in parenthesis. Significance: * p<0.1, ** p<0.05, *** p<0.01.

C Including the winter season

All of our estimations are from a sample that does not include the winter season. Sunshine is generally lacking during the winter season. Even with a clear sky, there may be no sunlight during showings, which in turn gives rise to noise. Moreover, the winter season is not a season in which people typically want to sell their houses. First, snow and ice will hide important characteristics and make units less accessible. Second, the general weather is perceived differently due to factors such as snow cover, snow depth, and icy roads and pavements. Third, the is a general problem with illiquidity in the winter season. Overall, including the winter season likely adds noise. Results from estimating the main model (1) when including the winter season are presented in Table C.1. The repeat sales results show that the overall (pooling) effect is smaller than previously found. This is consistent with adding more noise. Still, there is a significant winter effect, possibly related to cloud cover correlating with other winter-specific factors not controlled for, such as snow cover, snow depth, and ice on roads and pavements. By including a dummy for the winter season ($D_{wi.}$), the effect in the winter still is strong but not significant, and due to the additional noise in the cloud cover variable, the overall effect is still smaller and only significant at the 10% level. Here, within-season estimates are not significant. To deal with the added noise, we interact the dummy with the non-interacted cloud cover variable in (1) while leaving all other interactions unchanged. This makes the overall effect larger than in our main results, but because we do not allow for a different slope of the cloud cover interacted with hours of night during the winter, the in the spring and fall the estimate is smaller and significant at the 10% level. In the last column, we interact all cloud cover and weather variables, but not the de-seasoned controls, with the winter dummy, denoted by $\text{WxD}_{wi.}$. Allowing for different slopes for all weather variables and their interactions in the winter yields significant estimates with magnitudes similar to our main results in Table 2.

Table C.1: Winter results

	Cross-section				Unit FE			
		$D_{wi.}$	$CCxD_{wi.}$	$WxD_{wi.}$		$D_{wi.}$	$CCxD_{wi.}$	$WxD_{wi.}$
Marginal effects at the means, pooled and by season								
Pooled	-0.0059	-0.0056	-0.0152**	-0.0164**	-0.0149*	-0.0146*	-0.0249***	-0.0273***
	(0.0059)	(0.0060)	(0.0068)	(0.0069)	(0.0077)	(0.0077)	(0.0088)	(0.0090)
Spring	-0.0123**	-0.0106*	-0.0115*	-0.0110*	-0.0134*	-0.0123	-0.0131*	-0.0152**
	(0.0059)	(0.0060)	(0.0061)	(0.0061)	(0.0075)	(0.0078)	(0.0078)	(0.0079)
Summer	-0.0255***	-0.0235***	-0.0228***	-0.0228***	-0.0094	-0.0080	-0.0074	-0.0089
	(0.0083)	(0.0083)	(0.0083)	(0.0084)	(0.0106)	(0.0106)	(0.0106)	(0.0107)
Fall	0.0076	0.0087	0.0053	0.0075	-0.0183*	-0.0175	-0.0208*	-0.0232**
	(0.0076)	(0.0085)	(0.0088)	(0.0090)	(0.0098)	(0.0109)	(0.0113)	(0.0115)
Winter	0.0152*	0.0089	0.0135	0.0083	-0.0208*	-0.0246	-0.0212	-0.0189
	(0.0087)	(0.0136)	(0.0140)	(0.0161)	(0.0114)	(0.0194)	(0.0199)	(0.0223)
N	63,752	63,752	63,752	63,752	17,718	17,718	17,718	17,718
Adj. R. sq.	0.8737	0.8738	0.8738	0.8738	0.9589	0.9589	0.9589	0.9589
Fixed effects	MYQ	MYQ	MYQ	MYQ	YQ	YQ	YQ	YQ
Zip code FE	✓	✓	✓	✓				

Notes: The table shows marginal effects at the means from regressing the log of the selling price on cloud cover when including the winter season, in which the selling price includes common debt if any. Columns are denoted with the extensions to the main specification (1): $D_{wi.}$ is a dummy for the winter season, $CCxD_{wi.}$ is the interaction between the cloud cover and the winter dummy (not interacted with the other cloud cover interactions), and $WxD_{wi.}$ represents the interactions between the all weather variables with the winter dummy. Weather variables are the moving averages of the four days prior to the sale. All variables are in natural logarithms. MYQ and YQ denote the municipality-by-year-by-quarter and year-by-quarter fixed effects, respectively. All regressions include the weather controls maximum temperature and wind speed and their seasonal monthly averages. All regressions control for age and mortgage rate. Cross-section regressions control for attributes: size (square meters), an indicator taking the value of one if the housing unit is an apartment, the interaction between size and the apartment indicator, an indicator taking the value of one if the housing unit is an owner-occupier unit, and an interaction between lot size (square meters) and a non-apartment indicator. Heteroskedasticity-robust standard errors are given in parenthesis. Significance: * $p<0.1$, ** $p<0.05$, *** $p<0.01$.

4 How directed is housing search?

How directed is housing search?*

Andreas Eidspjeld Eriksen†

November 11, 2024

Abstract

Despite housing market search behavior being likely influenced by factors such as list prices, home attributes, and location, search theory often assumes random search. This paper investigates the predictive power of list prices on buyer arrival at listed housing units using machine learning, with the purpose of providing empirical evidence of the importance of directed search for buyer arrival. The paper utilizes a rich and unique dataset of housing transactions that includes information about buyer arrival at different stages of the selling process. The findings indicate that the overall importance of the list price is moderate at best, with exceptions in the lower-priced housing unit segment, where directed search appears to be a determinant of buyer arrival.

Keywords: *Housing market; Directed search; Machine learning*
JEL classification: *R21; R31*

1 Introduction

When listing housing units on online marketplaces, the primary objective for achieving a desirable price is to attract attention from potential buyers. The first stage in attracting buyers is listing the units online. This is followed by arranging open houses (hereafter showings), which potential buyers may attend. After the showings, some of the showing participants may make offers. These stages constitute the arrival process. A universe of potential factors may play into how and whether potential buyers end up at each stage, and it is generally difficult ex-ante to determine the interest for a particular housing

*I am grateful to André Kallåk Anundsen, Aslak Bergersen, Cloé Garnache, Brett Lantz, Erling Røed Larsen, Plamen Nenov, Dag Einar Sommervoll, and seminar participants at Ambita/Boligmappa and Eiendomsverdi for helpful comments.

†School of Economics and Business, Norwegian University of Life Sciences and Housing Lab, Oslo Metropolitan University. E-mail: andrease@oslomet.no

unit. The universe of factors motivates the modeling of the buyer arrival process as the outcome of a random housing search process. Despite this, list prices are often assumed in the literature to be able to help buyers direct their search across housing markets.

This paper investigates whether buyer search flows can be considered as directed by assessing the predictive power of the list price on buyer arrival at listed housing units. I apply machine learning techniques to model buyer arrival both with and without the list price as a predictor. The models include variables that may affect buyer arrival and also correlate with list prices. These variables are important for isolating the predictive power of the independent information provided by list prices: from a directed search perspective, they should be informative beyond their capacity to nest relevant information about the listed housing units. These predictors include hedonic attributes, market factors including tightness, weather variables, and previous arrival stages. The results from comparing out-of-sample predictions suggest that list prices have limited importance on average for explaining buyer arrival. However, this hides important heterogeneous effects: in the lower list price segment, increasing the list price is associated with a reduction in buyer arrival. This suggests that random search may be a valid assumption for some types of housing, while there are specific segments of the housing market where directed search appears to be an important determinant of buyer arrival.

Real-world housing market search is likely to be directed by list prices and partially segmented by hedonic attributes, simply because online marketplaces where housing units are listed offer search filters on these variables. When filtering on the list price, buyers can add more filters, such as size and the number of bedrooms, so that the list price may affect arrival in combination with other attributes. When modeling housing market search, the complexity of search behavior, together with the purpose of the theory model itself, can make it reasonable to assume that search is random. Assuming random search is the traditional setup from the Diamond-Mortensen-Pissarides (DMP) search framework for labor markets, initially applied to a housing market setting by Wheaton (1990). Yet, there is a strand of literature that allows search to be directed by list prices. There is also an emerging literature on segmented housing search that allows buyers to focus their search into different market segments (Piazzesi et al., 2020). Ultimately, this paper seeks to provide empirical insights into whether it is more plausible to use a setup based on random or directed search, by evaluating how important list prices are for predicting buyer arrival.

In directed search models, list prices help potential buyers focus their search. As described by Wright et al. (2021), directed search is part of the *environment*, or setting, of theory models, meaning it is part of the technology. Directed search differs from segmented search where buyers search within segments of the housing market, as proposed and documented by Piazzesi et al. (2020). In real life, it seems implausible that buyers search completely at random, thus matching with sellers at random, either in the market or within sub-markets. This is also inefficient as random search models typically do not lead to efficient resource allocation (Engelhardt & Rupert, 2017; Moen, 1997). With the increasing influence of the online marketplaces in housing markets, which allow for filtering on home attributes and list prices, taking the directed search approach seems ever more plausible today. An example of how directed search can be implemented is Carrillo (2012), who permits potential buyers with (conditional) reservation prices above the list price to participate in the showing. Although not considering the choice of which ads to inspect, the model simplifies the process of buying a house when buyers can avoid spending time visiting units that are potentially bad matches for them. Other studies have also used directed search models that are able to explain important aspects of housing markets, such as Rekkas et al. (2022) who provide a model consistent with residual sales price dispersion and list price stickiness. Albrecht et al. (2016) and Han and Strange (2016) provide models consistent with the fact that housing units sell for less, at, and above the list price. Furthermore, Han and Strange (2016) provide empirical results of the relationship between the number of bidders and list prices, finding that increasing the list price is associated with fewer bidders. In total, using directed search can be crucial for explaining certain aspects of different markets such as housing markets.[1]

In general, arrival happens in the stipulated four stages: (i) the inspection of online listings which generates advertisement (ad) *clicks*, (ii) attending *showings*, (iii) registering *interest* in receiving updates on offers, and (iv) making offers to the seller. Buyers making offers are called *bidders*. In each stage, potential buyers acquire more information about the housing unit, including the common value component, and decide whether to attend the next stage. Thus, the process filters out potential buyers with lower match quality, leaving in the last stage those willing to make a binding offer.

To assess whether search can be considered as directed, I compare the predictive power of predictions produced by two machine learning models: one that includes the list price as a predictor, and one that leaves the list price out. In directed search models, list prices provide buyers with independent information about the likely price the housing unit will be sold at due to seller-specific heterogeneity. That is, for identical houses, list prices signal seller heterogeneity.[2] In addition, list prices may also be set strategically by the sellers,

[1]There are also studies providing empirical evidence of the role of the list price for time-on-market, for instance Andersen et al. (2022), Anglin et al. (2003), Anundsen et al. (2022), Guren (2018), Haurin et al. (2010), and Yavas and Yang (1995). These studies typically show that higher list prices are related to longer time-on-market, which is argued to be caused by sellers being reluctant to sell below the listed prices. This is not necessarily consistent with directed search.

[2]Specifically, sellers may use list prices to signal buyers about their types Albrecht et al. (2016). They can also use list prices as commitment devices (Y. Chen & Rosenthal, 1996), serving as ceilings for selling

adding noise to the potential pure directed search effect of list prices. List prices may also provide independent information about housing-specific heterogeneity, unobserved by the econometrician, which can help segment buyer search. From these perspectives, the list price should contain information unobserved to the econometrician, thus matter for buyer arrival. How much the list price matters for buyer arrival from other predictors will quantify the degree of directed search, or of segmentation on unobserved attributes, in the housing market.

To test the predictive impact of list prices, I predict buyer arrival in each stage of the search process using data from the Oslo housing market in the period 2018-2023. The variables used as predictors include hedonic attributes of the housing units, location including neighborhood characteristics and distance to amenities, time of sale, market factors including market tightness, and the weather. The previous arrival stages are also used as predictors, for instance, when modeling the number of showing participants, ad clicks is used as a predictor, and when modeling the number of bidders all other arrival stages are used as predictors. The prediction problem is potentially complex, with a universe of variables, functional forms, and interactions that could affect arrival. Therefore, I use gradient boosted trees as implemented through eXtreme Gradient Boosting (hereafter XGBoost) to model buyer arrival (T. Chen & Guestrin, 2016). Predictions are made on a holdout sample. The difference in predictive performance when including and leaving out the list price is the main indicator for the importance of the list price as a predictor. Moreover, the models themselves can tell how the list price affects predictions across the list price distribution. These patterns are extracted using SHapley Additive exPlanations (SHAP) values (Lundberg & Lee, 2017), which reveal how each prediction is affected by each predictor similarly to coefficient estimates in linear regression.

The results suggest that the list price is at best moderately important for predicting ad clicks. The list price is even less important for showing participants, interested people, and the number of bidders. However, the average predictive power hides important heterogeneous effects. Investigating the SHAP values reveals nuances in how list prices affect buyer arrival. In the lower segment of list prices, being prices below 4-4.5 millions Norwegian Kroner (approximately 400,000-450,000 USD), they have greater importance in predicting ad clicks. In this segment, increasing the list price will decrease the number of ad clicks. The marginal effect of list prices above this price threshold is positive, meaning, the higher the list price the more ad clicks. The negative marginal effect for lower list prices is consistent with directed search. The pattern for higher priced housing units suggests that these list prices carry more information about unobserved housing unit quality: higher priced units may be more heterogeneous than lower priced units.

prices, or as partial commitments (Han & Strange, 2016).

Lower priced units show the same pattern for the number of bidders. However, here the marginal effect is monotonic: the higher the list price, the less bidders should arrive. This monotonic pattern is not evident in earlier arrival stages, which is consistent with the fact that making an offer constitutes a real commitment because offers are non-retractable. The cost of arriving in the earlier stages is relatively negligible compared to making offers. Yet, even with a meaningful impact of the list price, the magnitude of the marginal effect is limited, so economically, the list price is not the most important determinant of bidder arrival. In total, directed search appears to be an important determinant for buyer arrival only in certain market segments, particularly in the segment for cheaper housing units.

The findings contribute to the existing housing market search literature (such as Albrecht et al. (2016), Anenberg and Bayer (2020), Carrillo (2012), Y. Chen and Rosenthal (1996), Genesove and Han (2012), Guren (2018), Han and Strange (2016), Head et al. (2014, 2016, 2018), Jiang et al. (2024), Moen et al. (2021), Ngai and Sheedy (2020), Novy-Marx (2009), Piazzesi et al. (2020), Rekkas et al. (2022), and Wheaton (1990)) by providing reduced-form empirical results of search behavior from a real-world housing market. The paper adds to the literature that empirically explores search such as Banfi and Villena-Roldan (2019), Faberman and Menzio (2018), and Marinescu and Wolthoff (2020), with a housing market specific study of online advertisement being Barnwell and Fournel (2022). However, this paper is closer to Engelhardt and Rupert (2017) who compare competitive (directed) versus random search models, using calibration methods in a labor market setting. They rejecting mostly the competitive search approach.

2 Institutional setting

When selling a housing unit in Norway, the seller hires a realtor. The realtor serves as a mediator and is required to provide both the seller and potential buyers with advice.[3] After the realtor is hired, she gathers all relevant information about the housing unit, puts together the documentation and makes a prospectus and an advertisement. Advertisements are usually only listed online at the internet marketplace Finn.no, which holds monopoly power for housing market listings.[4] Publishing an ad on Finn.no constitutes listing the housing unit on the market, and showing dates are usually announced either at the time of listing or a few days later. Showings are arranged by the realtor. After the showings, the realtor arranges the sale by keeping track of incoming offers, informs those on the list of interested people about new offers, and keeps in contact with all ac-

[3] By law, realtors must act in accordance with "care for the interest of both parties" (Eiendomsmeglingsloven, 2007, §6-3).

[4] A new competitor to Finn.no, hjem.no, was announced in March 2024 which has gained some traction. Finn.no had monopoly power in the sample period.

tive parties.[5] When multiple buyers submit offers the sale process can resemble an open ascending bid auction with buyers bidding against each other by submitting higher offers. For this reason, and given the frequent occurrence of these bidding wars in Norway, I refer to the sale as an *auction* and the offers as *bids*. Bids received by the realtor are legally binding, meaning they are considered as offers which cannot be withdrawn. If a bid is accepted by the seller, the winner is legally required to pay the agreed upon price which is the amount of the accepted bid.

Figure 1: Buyer arrival sequence

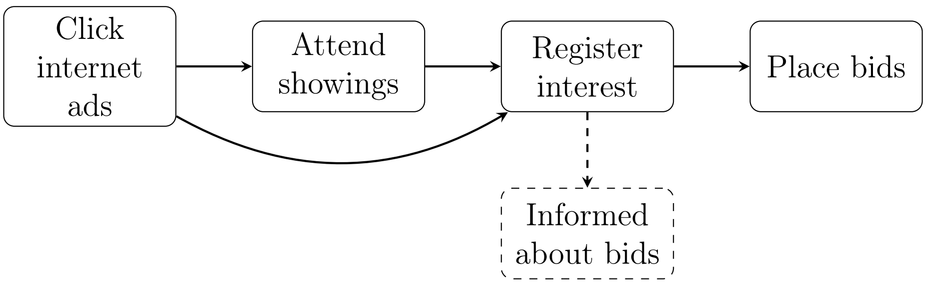

To summarize, buyers use Finn.no to search for housing units, adding filters to find relevant listings, which in turn accumulates ad clicks for the housing unit. Some buyers with high match quality will find it worthwhile visiting the housing unit at the showings, but some might also contact the realtor for arranging a private showing.[6] At the showings, buyers can sign up to being informed about incoming offers, but this registration can also be done by contacting the realtor directly at any point in time. After the showings are finished, buyers place offers, bidding against each other if more than one buyer participate in placing offers. This process constitute the sequence of buyer arrival to housing units, all of which are outcomes of interest. The arrival process is visualized in Figure 1.

3 Data

The data comprise transactions from DNB Eiendom, the second largest real estate agency in Norway. The sample covers on-market transactions in Oslo in the period September 2018 to May 2023. Relevant variables in the data are the total ad clicks, the number of showing participants, the number of interested people, within-auction bidder ID, the address, the list price, the selling price, the common debt, housing characteristics, the time of listing, the time of sale, and variables related to the bids themselves. All prices add the common debt if any. The total number of ad clicks are divided by the time-on-

[5]This includes one-on-one contact with everyone on the list and the seller.

[6]Private showings rarely occur and I disregard these in the analysis. About 1 percent of transactions in the raw sample have private showings.

market, meaning the number of days between the listing and the sale, and it is divided by 1,000. The sample is cleaned and restricted to contain well-behaved transactions. An in-depth explanation of the steps taken to obtain the final sample is provided in Appendix A. Summary statistics of the main variables are provided in Table 1.

Table 1: Summary statistics

	Mean	SD	Min	1st quartile	Median	3rd quartile	Max
Ad clicks per day (in 1,000)	0.40	0.21	0.13	0.25	0.35	0.50	1.30
Showing participants	16.91	13.99	0.00	7.00	14.00	24.00	69.00
Interested people	15.04	10.46	1.00	7.00	12.00	20.00	57.00
Bidders	2.72	1.57	1.00	2.00	2.00	4.00	8.00
List price (MNOK)	4.53	1.62	2.49	3.42	4.03	5.09	11.54
List price/size (in 1,000)	74.34	21.87	26.05	57.90	73.97	88.82	207.69
Size (square meters)	65.46	28.03	13.00	47.00	62.00	77.00	299.00
Age	59.50	33.63	1.00	39.00	59.00	81.00	185.00
Bedrooms	1.73	0.94	0.00	1.00	2.00	2.00	7.00
Lot size (square meters)	16,454.65	24,217.23	21.10	1,353.05	6,006.00	22,708.00	765,591.00
Unit type: Detached house	0.01						
Unit type: Apartment	0.93						
Unit type: Row house	0.04						
Unit type: Semi-detached	0.01						
Ownership: Non-co-op	0.37						
Ownership: Co-op	0.57						
Ownership: Other	0.06						

Notes: The table presents summary statistics of the full sample. The sample consists of well-behaved transactions in Oslo, with a total of 8,263 observations. The cleaning steps in Table A.1 reports that the final number of observations is 8,271, but there are some missing values in the sample thus the discrepancy in observations. When missing, the lot size is mean imputed conditional on housing unit type. The ad clicks per day is the total number of ad clicks divided by time-on-market, and it is reported in 1,000 clicks. The list price is reported in millions of Norwegian Kroner (MNOK), and include common debt if any. The size refers to the living area in square meters.

In order to add more *signal* to the models, I collect a large set of additional information about the housing units, the area in which they are located, market factors, and the weather. I find walking distance to parks, primary schools, kindergartens, public housing addresses, tram stops, subway stations, and points of interest along the coast line. For grocery stores and driving entrances to the main highway that goes through Oslo (Ring 3 and E6), I calculate travel distance by car. I find the distances, as the crow flies, to the forest areas Nordmarka and Lillomarka, the coastal line, and areas with high degree of noise pollution from traffic, trains, subway, and trams. Neighborhood characteristics are collected from the City of Oslo on the *basic statistical unit* (BSU) level, which is a small geographical area. For each BSU, I calculate annual population density, share of population with different country backgrounds (first generation only), total housing units as a fraction of population, number of each housing unit type as a fraction of population, and the number of each housing unit type as a fraction of the total number of housing units. The market factors include the annual frequency of which housing units are transacted on the market in each BSU, and the buyers to sellers (B/S) ratio lagged one month is meant to proxy for the market tightness. Weather variables include temperatures, precipitation, snowfall, snow depth, and wind, all of which are included as moving averages between different key dates such as the date of listing and the date of the final showing. Appendix A goes into more detail about both the additional variables and the treatment of these in preparation for model estimations.

4 Method

4.1 Predictive performance and directed search

Using the list price as a predictor for buyer arrival to assess whether their search is directed requires two sets of models: one that includes the list price and variables derived from it as predictors, and one that does not. The predictive impact of the list price, inferred by comparing performance, should indicate how important the list price is for buyer arrival.

While this approach makes it possible to measure the list price importance, there are two perspectives to be considered for explaining why the list price should matter. The directed search perspective is that list prices give signals about the seller heterogeneity. Sellers reveal their information about themselves by setting list prices that deviates from other sellers, for instance, signaling the urgency or willingness to sell at a lower price. As a predictor, the list price should help the buyers distinguish the sellers, making it an important predictor for buyer arrival flows. Although not consistent with directed search theory, strategic price setting is found to affect the probability to sell.[7] Some sellers may try to follow such a strategy, which leads to noise when distinguishing out the heterogeneity-signaling effect from directed search.

The other perspective relates to unobserved heterogeneity of housing units. The list price should carry important information about buyer arrival that is hard to observe, for instance, the need for renovation. Two housing units appearing to be identical in the data may have widely different qualities, and this quality discrepancy may be evident from the list prices.

Both of these perspectives provide a reason for why the list price should matter for buyer arrival. Modeling buyer arrival while accounting for a large set of predictors should help limit the impact of unobserved heterogeneity when comparing predictions of buyer arrival. To that end, some heterogeneity is harder to observe, and some of the unobserved heterogeneity may be functional forms and interactions. The more effort put into collecting information about the housing units should increase the likelihood of quantifying the

[7] See e.g. Anundsen et al. (2022) and Yavas and Yang (1995).

degree of directed search, and limit the magnitude of the list price capturing unobserved quality. Therefore, provided the wide set of variables described in Section 3, the difference in predictive impact on buyer arrival between including and leaving out the list price should provide evidence about the degree of directed search.

4.2 XGBoost

The choice of method takes account of the two-fold purpose this must serve, namely to be able to train a model that can account for the complexity of the prediction problem, and to some extent be interpretable. Therefore, I use gradient boosted trees as implemented by the method XGBoost, see T. Chen and Guestrin (2016). XGBoost satisfies both of these purposes: it is highly regarded as one of the most powerful methods for machine learning purposes (Lantz, 2023, p. 629), and it is possible to extract how the model treats the list price when generating predictions through explanatory tools such as SHAP values (Lundberg & Lee, 2017). SHAP values are provided and interpreted at the level of the unit of observation, and they will provide insight into how the list price affects model predictions across the list price distribution. Implying, even if the list price does not seem to provide predictive power on average, this could differ across price segments.

In short, XGBoost builds decision trees in sequence (T. Chen & Guestrin, 2016). At each iteration, the tree that optimizes (minimizes) the objective function is the next tree to be built. This corresponds to a sophisticated data-driven *model selection* procedure. The objective function consists of two parts, namely a loss function and a regularization term. The loss function is chosen by the user based on the prediction task at hand, and the regularization penalizes model complexity, which is also chosen by the user. There are several loss functions available for numeric prediction in XGBoost.[8] These provide flexibility with respect to the outcome variable distribution, and include among others Gamma, Poisson and Tweedie loss functions, all of which use the negative log-likelihood loss.[9] The default loss function is the squared error.[10]

4.3 Performance metrics and impact of list price

To measure how well the models perform out of sample, I calculate the mean absolute error (MAE) and the pseudo R^2, which squares the Pearson correlation between observations and predictions. My preferred performance metric is the MAE because it is robust to cases when the prediction error distribution deviates from the normal distribution (Hodson,

[8] XGBoost also support user-defined loss functions.

[9] XGBoost supports the Tweedie distribution family with a variance power in $(1, 2)$, meaning that it is a compound Poisson-gamma distribution.

[10] For more information, see e.g. the XGBoost website: https://xgboost.readthedocs.io/en/stable/.

2022), and because it is symmetric in dispersion (Steurer et al., 2021).[11] The pseudo R^2 is used as an universal metric of goodness-of-fit for comparison across outcomes. I also calculate the root mean squared error (RMSE) and the mean bias error (MBE). RMSE is used as one of the loss functions, making it inappropriate to measure performance when allowing for other loss functions. RMSE and MBE are reported in the Appendix.

To evaluate the performance impact of the list price, models of the outcomes are trained twice: for each outcome the predictions from including and leaving out the list price are compared based on the prediction MAE. I measure improvement in predictive power by assessing the change in MAE using these measures:

$$\Delta\text{MAE} = \text{MAE}_{\text{list}} - \text{MAE}_{\text{no list}} \tag{1}$$

$$\Delta\%\text{MAE}/s = \frac{\Delta\text{MAE}}{s_y} \times 100. \tag{2}$$

Meaning, to make sense of the change in MAE, the difference ΔMAE is compared to the sample standard deviation (SD) s_y of the outcome variable y. Hence, the potential improvement is compared relative to the dispersion level of the respective outcome.

4.4 Considerations about estimation

When estimating models for predicting the outcomes, there are several considerations to take into account. First, for benchmark predictions, I consider both ordinary least squares (OLS) and generalized linear methods (GLM) of different distributions. When using standard regression methods, the chosen methods should conform with the distributional properties of the outcome variable.[12] I choose the GLM variants based on the outcome distributions, as presented in Figure 2. The daily ad clicks in thousands, which I consider as being continuously distributed, seems to be close to Gamma distributed. The discrete showing participants is zero inflated: the distribution has a clear bunching at zero. When truncating at zero, the negative binomial (NB) distribution fits well. The number of interested people fits also well with the NB. Finally, the Poisson and the NB distributions both seem to be good fits for the bidders, indicating equi-dispersion. These insights are also used to choose between loss functions when training XGBoost models.

The second consideration is about market effects and exogenous factors that may at least appear stochastic for the agents in the market. The key market factor to account for is the *market tightness* measured as the buyers to sellers ratio (B/S). Market tightness

[11] Symmetric in dispersion means that both tails of the distribution are treated equally.

[12] Typically, using ordinary least squares to predict an outcome with bounded support can result in predictions outside of the bounds.

Figure 2: Distributions of outcomes

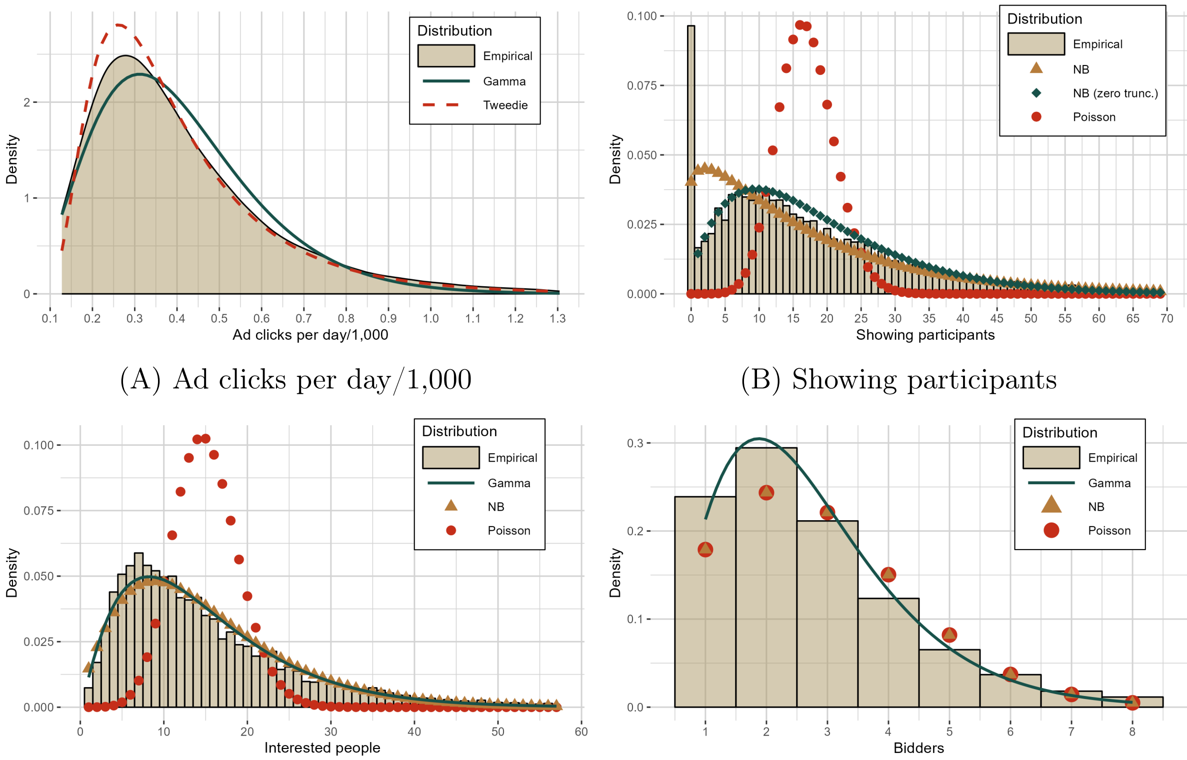

(A) Ad clicks per day/1,000

(B) Showing participants

(C) Interested people

(D) Bidders

Notes: The figure presents distributions of the four outcome variables of interest, with density functions of theoretical distributions fitted with maximum likelihood.

is a key component in any search model through the use of a matching function.[13]

Another market factor that can be important for arrival is the frequency at which units in the BSU are sold, both in total and for housing unit types such as apartments and detached houses. Other perceived stochastic factors are the effects the weather have on arrival. Weather can make a housing unit look better or worse (Gourley, 2021); it can be important for opportunity costs of search versus other activities; and in cases of bad weather, or much snow, it can constrain mobility thus showing attendance while inflating ad clicks since people are stuck at home. Both tightness and weather can account for seasonality.

Thirdly, the reduced form approach to model search behavior should take account of the buyer arrival process in Figure 1. The reason being that these variables will contribute to isolating the directed search effect of the list price. Modeling the outcomes could potentially be harder further along the arrival process, because important information about why potential buyers ends up dropping out between stages are lacking from the data. For instance, if a potential buyer decides to go to a showing after seeing the ad,

[13]See, e.g., Jiang et al. (2024) and Novy-Marx (2009) for applications of time-varying market tightness in housing search.

data about the updated information set after attending the showing is lacking. Although not a part of the searchable set of variables, the earlier stages of arrival nest important variable subsets for the later stages, thus these variables are included.

Fourth, the sequence of the arrival process restricts the set of variables that can be used to predict outcomes. In addition to including arrival information from the previous stages, modeling for instance ad clicks should not include information about the weather the days prior to the sale.

4.5 Tuning hyper-parameters

To provide stable models and avoid over-fitting, I provide different combinations (a grid) of hyper-parameters for XGBoost to tune on, each being evaluated by a conventional 10-fold cross-validation. In general, a k-fold cross-validation trains the model on k-1 folds and keeps the remaining fold for testing the model, a process repeated k times. Moreover, for each combination of hyper-parameters, the number of boosting iterations is chosen by the implementation of early stopping which I set to 50 iterations. Early stopping makes sure that the model stops building new trees when the test fold predictions does not improve after 50 iterations. This process is repeated for the grid of pre-specified hyper-parameters. Table A.14 in the Appendix provides information about the grids.

The cross-validation requires a metric for evaluating performance. When deviating from the default squared error loss function to Gamma, Poisson or Tweedie loss functions, the evaluation metric is the negative log-likelihood of the respective distribution.[14] This will make the methods optimize on different metrics, which is an advantage since the preferred metric is the MAE. Meaning, across the different cross-validated XGBoost models, I report results from the model that performs the best based on MAE.

5 Results

5.1 Benchmark

Standard regression methods are used as benchmarks, consisting of OLS, Gamma GLM, Poisson GLM, and negative binomial GLM. The benchmarks use the list price, the list price per square meter, and a set of typical hedonic variables consisting of year-by-month fixed effects, 3 digit zip code fixed effects, ownership type, housing unit type, and housing unit attributes which include living area size and age. All numeric variables, including

[14]Specifically for the Tweedie loss function, it requires the variance power of the distribution, which I find by fitting a theoretical Tweedie to the outcome with maximum likelihood. If the power exceeds the supported variance power limits (the supported variance power is in $(1, 2)$), I choose the power at the limit with a small correction to ensure it is inside the supported interval.

the list price and the list price per square meter, are included taking the power of one to three. The predictors are chosen to be representative of typical hedonic variables used in sales price prediction, also being variables that are available for filtering when browsing listings online. When leaving out the list price, this means leaving out the list price per square meter and all the functional forms. No other interactions are considered here.[15]

Table 2: Benchmark results

	No list		List		Improvement	
	MAE	Δ_I%MAE	MAE	Δ_I%MAE	ΔMAE	Δ%MAE/s
Ad clicks per day						
OLS	0.132	-19.551	0.129	-21.323	-0.003	-1.386
Gamma	0.131	-20.260	0.130	-20.636	-0.001	-0.294
Showing participants						
OLS	8.620	-20.373	8.629	-20.295	0.008	0.060
NB	8.636	-20.232	8.604	-20.523	-0.031	-0.225
Interested						
OLS	6.423	-21.166	6.408	-21.342	-0.014	-0.137
Gamma	6.364	-21.883	6.328	-22.326	-0.036	-0.345
NB	6.357	-21.966	6.318	-22.446	-0.039	-0.373
Bidders						
OLS	1.221	-2.813	1.195	-4.846	-0.026	-1.626
Gamma	1.221	-2.748	1.192	-5.052	-0.029	-1.842
Poisson	1.222	-2.734	1.192	-5.103	-0.030	-1.894

Notes: The table presents results from using different regression methods for predicting the four outcome variables, with the *Ad clicks per day* being in 1,000 clicks. OLS, Gamma, NB and Poisson all refer to regression methods, with *NB* referring to negative binomial regression. Columns are grouped by *No list* and *List* referring to leaving out and including list price predictors, respectively. The columns denoted Δ_I%MAE give the increased predictive performance from an *intercept model* in percentage terms, meaning, OLS with just an intercept, which is equal to using the average of the outcome in the training sample as predictions. The column ΔMAE is the difference between MAE when including list price and MAE when not including list price, and Δ%MAE/s represents this difference relative to the sample standard deviation of the corresponding outcome in percentage terms. When included, list price is included both as nominal list price and as nominal list price per square meter (size of housing unit), also taking the squared and cubic functional forms of these. All regressions include the variables living area size, age, also taking the square and cubic terms of these, ownership type, housing unit type, 3 digit zip codes fixed effects, and year-by-month fixed effects.

The results in Table 2 show that the standard OLS and the more appropriate methods perform similarly. All outcomes are modeled with a baseline intercept model: an OLS regression with an intercept, resulting in predictions being the unconditional average outcome in the train sample. A reduction in MAE from the intercept model indicates

[15] When including year-by-zip-code dummies in the OLS models, this yielded very similar results.

an improved performance. The improvement in predictive power of the model without the list price can be inferred from the *No list* column group. Column Δ_I%MAE gives the change in MAE from the intercept model, showing that the hedonic predictors have some predictive power. For instance, using OLS to predict ad clicks per day (divided by 1,000) results in a MAE of 0.132, meaning that the model misses on average by 132 clicks. This corresponds to a reduction in MAE by almost 20 percent compared to the intercept model. The *List* columns provide prediction results from including the list price predictors in addition to the hedonic predictors. Here, predicting ad clicks per day using OLS results in a MAE of 0.129. The two last columns provide the improvement from including the list price predictors, meaning from *No list* to *List*. The improvement for the OLS model of ad clicks per day is not particularly high: the improvement amounts to missing by 3 ad clicks less. This reduction amounts to 1.386 percent of the sample SD.

The most improvement in MAE is when predicting the number of bidders, which is the Poisson model that ends up with a reduction in MAE of 1.894 percent of the sample SD. What is striking is how little predictive power the hedonic variables have for this outcome, as the MAE only reduces by 2.734 percent from the baseline intercept model. The Poisson model with list price gets a MAE of 1.192 which is close to 40 percent of the average of 2.72 and 76 percent of the sample SD at 1.57. The reduction from the baseline model is just 5.103 percent.

Adding to this, Table B.1 in the Appendix reports more metrics of the benchmark predictions, also including results from the baseline intercept models. The R^2 is very low for the number of bidders as it never exceeds 0.07 even when including the list price. Across all outcomes, the R^2 does not exceed 0.35. Overall, there seems to be little predictive power of the list price predictors for any of the outcomes.

5.2 Predictive performance impact using machine learning

Table 3 shows the performance metrics MAE and R^2 from predicting the four outcomes when leaving out and when including the list price. The list price is included as the nominal list price. All metrics can be found in Table B.2 in the Appendix.

Starting with the ad clicks, the model without the list price achieves an MAE of 0.117, which is better than the benchmark that achieved 0.131 in Table 2. Including the list price increased performance by reducing MAE by 0.003, which is equivalent to 3 ad clicks per day listed on the market. Recalling that the sample SD of ad clicks per day is 0.21, this increased performance is small as the model still misses on average by 0.114 or 114 clicks. The relative magnitude of this is about 54 percent of the sample SD, and the improvement amounts to 1.433 percent of the sample SD.

The MAE for showing participants is 7.264 without the list price and 7.245 with the

Table 3: ML results

	No list		List		Improvement	
	MAE	R^2	MAE	R^2	ΔMAE	Δ%MAE/s
Ad clicks per day	0.117	0.433	0.114	0.454	-0.003	-1.433
Showing Participants	7.264	0.494	7.245	0.498	-0.019	-0.136
Interested	4.079	0.713	4.075	0.715	-0.004	-0.037
Bidders	0.938	0.390	0.934	0.398	-0.004	-0.277

Notes: The table presents the main results of predicting the outcome variables. Ad clicks per day is in 1,000 clicks. The column ΔMAE is the difference between MAE when including list price and MAE when not including list price, and Δ%MAE/s is this difference relative to the sample standard deviation of the corresponding outcome in percentage terms. The loss functions used in the XGBoost grid searches are the default *squared error* (all outcomes), the *Gamma* (ad clicks, interested, bidders), the *Poisson* (showing participants, interested, bidders), and the *Tweedie* (ad clicks). The gradient and the hessian of the loss function, which is part of the objective function, are used for optimizing the objective function in each iteration. The loss function for *Gamma*, *Poisson*, and *Tweedie* is the negative log-likelihood of the respective distribution.

list price. The latter is equivalent to missing on average by about 52 percent of the sample SD. The rounded off improvement of including the list price is merely 0.019 which is about 0.1 percent of the sample SD, a very small improvement.

The predictions of the interested people show a similar pattern to the predictions of showing participants, with a very small increase in improvement when including the list price. The decrease in MAE is found to be 0.004 which is 0.037 percent of the sample SD. However, even without the list price the predictions are the best across all outcomes, with a MAE of 4.075. This translates to that the model misses on average by just 39 percent of the sample SD. This strong performance may seem strange when comparing with the other outcomes, but recall that each stage of the model uses the previous stages as predictors. Thus, the showing participants are included when modeling the interested people, and many of the showing participants sign up at the showings as being interested in participating in the auction.

Finally, the impact is not as strong for the number of bidders. When including the list price the MAE is 0.934, which is a decrease of 0.004 amounting to 0.277 percent of the sample SD. Although this translates to missing on average by 0.004 bidders less, it still achieves the second highest improvement across all outcomes when taking account for the outcomes dispersion.

It is surprising that the list price is not more important for the number of bidders. Placing a bid implies a big financial commitment, while the previous stages does not entail such commitments. Potential buyers may use the list price to direct their search in the first stage, but conditional on everything else, this effect disappears later in the buyer arrival sequence. In the later stages, the ad clicks may carry the list price effect. Yet, the magnitude of the importance is surprising: at most the magnitude of the performance improvement comprise only 1.4 percent of the sample SD. Meaning, the impact of including the list price, when already including a larger selection of predictors, seems to be modest at best. This leads to the question about how important the list price is when augmenting the model by leaving out the previous arrival stages. By doing so, the list price signal carried by earlier arrival can be left out to isolate the impact of the list price. The degree of list price nesting by the arrival stages and the impact of predictive performance is investigated in Subsection 5.4.

5.3 Marginal list price importance with SHAP

Although the results in Table 3 indicate that the difference in average predictive performance is at best moderate, this does not reveal how the models use the list price: there could be heterogeneity in the importance of the list price. This is investigated by calculating SHAP values using the full sample of observations. SHAP values are calculated at the unit of observation level, with each value providing information about how the list price of that observation contributes to the predicted value. SHAP values can be used to assess the overall importance of a predictor based on how the model utilizes it. When comparing observations of a predictor with its SHAP values, the parallel to this relationship is coefficient estimates from OLS: the estimate is used to weight the observed list price, which is added to the predicted value. Taking the mean absolute value of SHAP values for a given predictor will reveal the average importance of that variable for predictions, resembling a scale-adjusted OLS coefficient estimate.

I rank the mean absolute SHAP values of all predictors to get a sense of the importance of the list price relative to other predictors. The rank is 3 for ad clicks per day, 10 for showing participants, 16 for interested people, and 3 for bidders. Meaning, the list price is a relatively important predictor in the ad clicks and bidders models. Still, there are other predictors that are similarly or even more important. SHAP summary plots ordered by the most important predictors are provided in Figures B.1, B.2, B.3 and B.4 in the Appendix.

Dependence plots of the list price are presented in Figure 3, showing how the list price contributes to producing the predictions of the respective outcome.[16] The SHAP values for ad clicks show that list prices below 3.2-3.4 million NOK (roughly 320,000-340,000 USD in early 2023 exchange rates) positively impact predictions. Starting at the low list prices, the impact on predictions decreases as the list prices increase, but only to a certain

[16]The dependencies uncovered by the models are not readily seen in the raw data. See Figure B.5 in the Appendix.

Figure 3: SHAP dependence of the list price

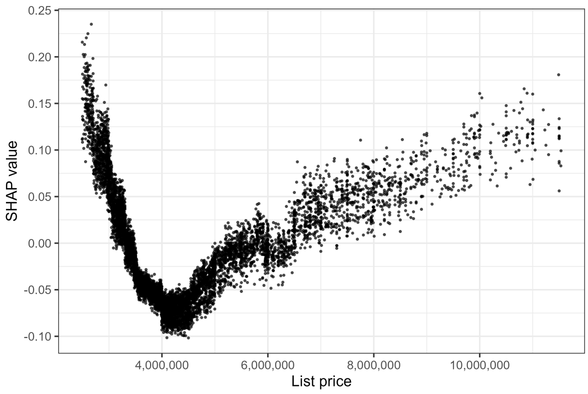

(A) Ad clicks per day/1,000

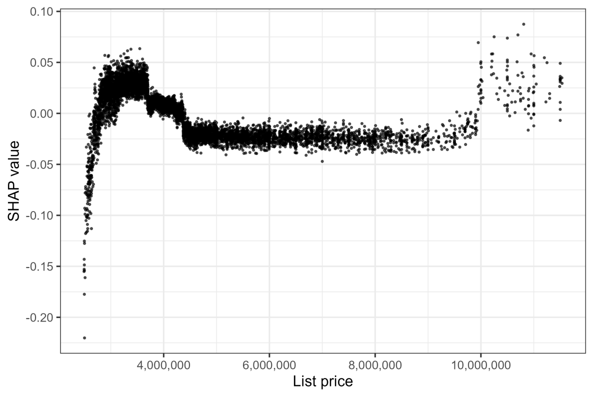

(B) Showing participants

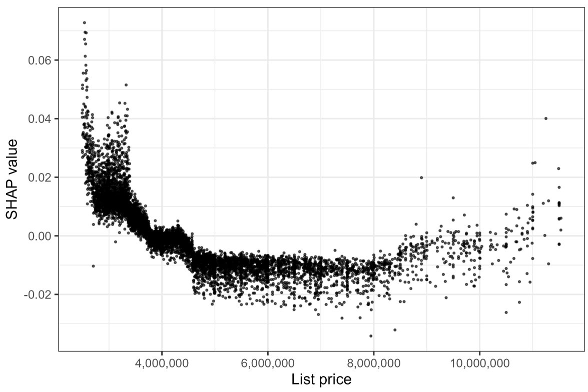

(C) Interested people

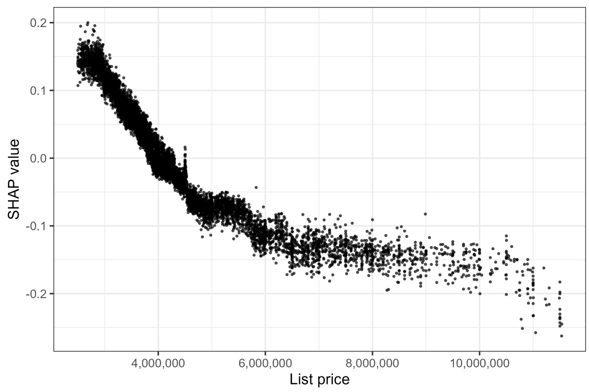

(D) Bidders

Notes: The figure presents SHAP dependence plots of the list price for the four outcome variables of interest. Positive SHAP values means that the predictor affects the outcome positively, and vice versa, while the magnitude tells how strong the predictor impacts the outcome.

point at around 4.1-4.3 million NOK. The pattern suggests a non-monotonic influence on ad click predictions.

List prices have relatively low impact on showing participants, and this is also evident from the individual SHAP values. The lowest list prices influence predictions negatively, while the local slope is positive and relatively steep compared to the other parts of the plot. This pattern is not evident for any of the other outcomes, but it may be related to higher variation in showing participants among the lowest priced units, which the model struggles with. There is a local maximum at 3.2-3.4 million NOK and for list prices above this point there is almost no impact on predictions from list prices. When considering the SHAP values of from the model of interested people, a similar pattern as found from the ad clicks model emerge: lower list prices have stronger impact on predictions, but for list prices above 4 million NOK there is almost no impact.

The lowest list prices also have stronger positive impact on bidder predictions, with a zero impact at around 4 million NOK. At 5 million NOK the slope flattens out but is still decreasing. Overall, the pattern looks monotonic. Considering the plots as a whole, there seems to be heterogeneity in the importance of list prices across the list price distribution. Lower list prices seem to be more important for arrival flows.

5.4 Nestedness of arrival stages

The results presented in Table 3 indicate nesting of the list price. This relates to the fact that the models include previous arrival stages as predictors that may nest the list price signal from previous stages to the current stage. For instance, when modeling showing participation, some of the signal of the list price may be carried, or nested, by the ad clicks. This might be the reason why list price becomes relatively unimportant while ad clicks are among the most important predictors. This may in turn downplay the importance of the list price in the later stages of arrival. Note that online ad clicks is the first stage, making it not use any arrival variables as predictors.

To assess this issue, I reiterated the modeling while leaving out the previous arrival stages. Thus, I effectively remove the source of potential nesting, which should result in a more isolated impact of the list price. The models still use all other information.

Table 4: ML results, without previous arrival stages

	No list		List		Improvement	
	MAE	R^2	MAE	R^2	ΔMAE	Δ%MAE/s
Showing Participants	8.406	0.332	8.340	0.351	-0.065	-0.468
Interested	6.125	0.368	6.081	0.382	-0.044	-0.421
Bidders	1.174	0.060	1.162	0.078	-0.012	-0.772

Notes: The table presents the results of predicting the outcome variables when leaving out previous arrival stages as predictors. The column ΔMAE is the difference between MAE when including list price and MAE when not including list price, and Δ%MAE/s is this difference relative to the sample standard deviation of the corresponding outcome in percentage terms. The loss functions used in the XGBoost grid searches are the default *squared error* (all outcomes), the *Gamma* (interested, bidders), and the *Poisson* (all outcomes). The gradient and the hessian of the loss function, which is part of the objective function, are used for optimizing the objective function in each iteration. The loss function for *Gamma* and *Poisson* is the negative log-likelihood of the respective distribution.

Table 4 reports the results from this exercise. These results reveal that including the list price has a larger impact than those reported in Table 3. This indicates that the previous stages of arrival do, to at least some extent, nest the list price signal. Although the impact is stronger now, it is still small relative to the sample moments. Showing participation achieves a reduction in MAE by 0.065 which is closer to 0.5 percent of the sample SD. The number of interested people also show this tendency, achieving a reduction of 0.044 which is closer to 0.4 percent of the sample SD. The absolute improvement in MAE on the number of bidders is found to be 0.012, which is larger than the improvement when using the previous arrival stages (0.004). These results suggest that the predictive power of the list price could be dependent on previous arrival, but the difference between the results in Table 3 is still very small.

It is worth noting the large drop in the general predictive power of bidders. Even when including the list price, the model is not able to achieve nearly the R^2 as when including the previous arrival stages: now it has an 0.078 compared to 0.398 when including arrival as predictors. It does not seem to be much relevant signal in the sample for the purpose of predicting the number of bidders.

The predictions in Table 4 suffer from leaving out arrival as predictors. While the potential issue with list price nesting is addressed by leaving arrival out, arrival variables should still be necessary to include as predictors to be able to assess the importance the list price. To assess the degree of which the list price is carried by arrival to the next stages of arrival, I take an econometric approach. Following the established sequence of arrival, ad clicks may nest list prices, showing participation may nest list prices and ad clicks, and so on. To untangle this pattern, I estimate a series of regressions to residualize (orthogonalize) and de-season each arrival stage. For simplicity, coefficients are given the same notation across specifications:

$$\mathrm{l}_{it} = \alpha + \delta_t + \epsilon_{it}^{l} \tag{3}$$

$$\mathrm{AC}_{it} = \alpha + \beta_1 \hat{\epsilon}_{it}^{l} + \delta_t + \epsilon_{it}^{AC} \tag{4}$$

$$\mathrm{SP}_{it} = \alpha + \beta_1 \hat{\epsilon}_{it}^{l} + \beta_2 \hat{\epsilon}_{it}^{AC} + \delta_t + \epsilon_{it}^{SP} \tag{5}$$

$$\mathrm{IP}_{it} = \alpha + \beta_1 \hat{\epsilon}_{it}^{l} + \beta_2 \hat{\epsilon}_{it}^{AC} + \beta_3 \hat{\epsilon}_{it}^{SP} + \delta_t + \epsilon_{it}^{IP} \tag{6}$$

in which l_{it} is the log of list price of unit i in year-by-month t. The rest of the abbreviations are AC for ad clicks per day, SP for showing participation, and IP for interested people. δ_t is the year-by-month fixed effects that are left in the sample after the data preparation procedure stipulated in Subsection A.2. $\hat{\epsilon}_{it}^{l}$ denotes the residual from estimating the relationship in equation (3). This procedure is intended to remove the explainable parts of each of these variables, starting with the list price then moving on to the arrival stages, in order to keep only the parts of these that does not carry the previous stages and list price.

Estimating the effect of both observed and unnested (previous) arrival should reveal the degree of importance of the list price and previous arrival stages in explaining arrival. I estimate two types of regression specifications for each arrival stage: (i) using one observed previous arrival at a time, and (ii) using all residualized previous arrival together. The results are reported in Table 5.

Table 5: Results from regressing arrival on unnested previous arrival

	Showing participants		Interested people			Bidders			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Ad clicks per day	23.658***	23.620***	19.759***		19.725***	2.404***			2.588***
	(0.808)	(0.798)	(0.583)		(0.464)	(0.110)			(0.094)
Showing participants				0.478***	0.426***		0.063***		0.058***
				(0.007)	(0.008)		(0.002)		(0.002)
Interested people								0.091***	0.059***
								(0.002)	(0.003)
Residualized arrival		x			x				x
Controls	x	x	x	x	x	x	x	x	x
Observations	6,198	6,198	6,199	6,199	6,199	6,199	6,199	6,199	6,199
R^2	0.431	0.432	0.461	0.629	0.650	0.145	0.286	0.319	0.357
Adjusted R^2	0.420	0.421	0.450	0.622	0.643	0.128	0.272	0.305	0.344

Notes: The table presents the regression results from regressing arrival outcomes on observed and unnested arrival. Unnested arrival variables are obtained from estimating models (3), (4), (5), and (6). Note that when estimating models (5) and (9) with observed arrival, R^2 is 0.648 and 0.354, respectively. The control variables include year-by-month fixed-effects, 3 digit zip codes, type of housing unit (apartment, detached house, semi-detached, or row house), age and size. Both age and size are also included taking the squared and cubic functional forms. The sample between dependent variables differs due to the stratified sampling used in the data preparation, being the reason for why there is one observation less for the showing participants. Heteroskedasticity-robust standard errors are provided in parenthesis.

There are no large differences in coefficient estimates for either of the arrival stages. The two first columns reveal that removing and de-seasoning ad clicks only yield a marginally smaller coefficient estimate, and it is still highly significant, indicating that ad clicks do not carry relevant information about the list price in explaining showing participants. The same pattern is observed for modeling the number of interested people. There are some minor changes in coefficient estimates when modeling the number of bidders, but the point estimate remains similar. Surprisingly, the R^2 is higher when using the residualized variables: when estimating relationships with residualized (observed) arrival, R^2 is 0.432 (0.431) for showing participants, 0.650 (0.648) for interested people, and 0.357 (0.354) for bidders. A potential reason for this may be that the residualization leaves out some noise in the explanatory variables. However, these differences are negligible. Overall, it does not look like the list price, or previous arrival, is substantially nested in the arrival variables.[17]

[17]Note that plots such as those in Figure 3 can shed light on the dependence between predictors. When conditioning on other predictors, the plots can reveal whether the SHAP spread at each point along the list price axis can be, to some extent, explained by other predictors. This is especially relevant for the bidders since they depend on three other arrival stages. The dependence plots for the model of bidders, that conditions on the previous arrival stages, partly supports this story: interactions with showing participants and interested people seem to capture some of the spread in the plot especially for list prices below 4.5 million NOK. However, the slopes does not seem to differ. The plots are presented in Figure B.6C in the Appendix.

5.5 Discussion

The prediction results show an overall tendency that list prices have small to moderate predictive power on arrival. The list price has the highest impact when predicting the ad clicks and the number of bidders, being among the most important predictors. A potential reason for the limited importance of the list price, which is assessed in subsection 5.4, is that the important predictive information (signal) of the list price may be partly nested in the previous arrival stages. Even if the list price seems to play a more important role as a predictor of arrival when leaving out previous arrival, the evidence suggests that this nested effect is very limited. A regression scheme to untangle the nested effect also supports this notion.

To assess whether the list price is important enough to attribute the results to being in favor of directed search, a pragmatic choice could be, for instance, to require that the reduction in MAE, in absolute terms, should be at least 10 percent of the sample standard deviation of the outcome. For the ad clicks and the number of bidders, this means that MAE should reduce by at least 0.021 and 0.157 respectively. None of the results are close to such a threshold, with the largest improvement relative to the sample standard deviation being 1.4 percent (ad clicks).

Yet, these aggregated average metrics hide interesting important patterns. Calculating SHAP values help extract how the models treat the list price as a predictor. Plotting these values reveal that the marginal effects on predictive importance are stronger for lower list prices. Although Figure 3 cannot be interpreted as causal relationships, they show how the models utilize list prices to produce the predictions. The ad clicks plot shows that increasing the list price tends to reduce the predicted ad clicks. Estimating simple linear regressions for list prices below 4 million NOK to obtain the slopes as seen in the plots reveals the following. An increase in the list price by 500,000 NOK is associated with a reduction of 76 (predicted) ad clicks per day. For the same part of the list price distribution, the same increase in the list price is associated with 0.008 more showing participants, 0.010 more interested people, and 0.058 more bidders.

The two most striking patterns are found for the ad clicks and bidders. Above the 4-4.5 million NOK threshold, increasing the list price seems to have a positive impact on ad clicks predictions, all else equal. Cheaper housing units in Oslo are typically smaller apartments, which are affordable for first-time buyers. During the sample period, there has been a high demand and a lack of supply[18] in Oslo which has resulted in increasing prices, and the downward slope could relate to market entrants being more budget constrained than current homeowners searching for new homes. Meaning, for cheaper housing units, list prices direct buyer search because an increase in the price provides buyers with

[18]Mainly due to insufficient construction of new housing units.

information about the potential selling price of the unit and the seller's the willingness to sell. The increasing slope for higher priced units may relate to more heterogeneity in this price segment. These units are most likely higher priced due to size and location, while higher price may also relate to qualities that are hard to observe in data. If so, list prices for higher priced units may carry more information about unobserved heterogeneity, ultimately leading to the upward slope.

For the number of bidders the pattern for lower priced units should be considered together with the showing participants and the interested people predictions. Offers from buyers to the sellers are binding, thus, the cost of clicking on ads or even attending a showing contrasts with the real commitment of placing bids. Even with many people exhibiting interest for an expensive housing unit, many people are either unable or not willing to make such commitments by placing bids. In contrast to this notion, the list price has an overall small predictive impact on the number of bidders. The small magnitudes of the SHAP values for showing participants and interested people suggest that list prices are less relevant for these outcomes across the list price distribution. For the lower priced housing units, the marginal effect of the list price on ad clicks and number of bidders is consistent with directed search, while directed search seems to be less of a determinant for arrival for the higher priced market segment.

6 Conclusion

Housing market search models can be separated into two categories, namely models that assume buyers are searching randomly and models that allow buyers' search to be directed by list prices. This paper tests the directed search approach by assessing how important the list price is for predicting buyer arrival flows. Buyer arrival to listed housing units happens in four stages: they click online ads, participate in public showings, register to participate in the auction (as interested), and place bids. To assess the importance of the list price, I predict arrival to housing units in the Oslo housing market over the period 2018-2023 using machine learning techniques. Under the assumption of random search, the list price should not have any predictive power in any of the four stages of arrival, while it should have some degree of predictive power if search can be regarded as directed.

Each stage of arrival is modeled using a large set of predictors including hedonic attributes, location including neighborhood characteristics and distance to amenities, time of sale, market factors including market tightness (buyers/sellers ratio), and the weather. The results suggest that the list price has a limited and potentially negligible predictive power on buyer arrival. It is more important for ad clicks than for showing participation, the number of interested people, and the number of bidders. Yet, there is heterogeneity across the list price distribution: the models for ad clicks and number of bidders show evidence consistent with directed search for cheaper units.

Future research may benefit from including even more predictors and observations, because this may help isolate the importance of directed search for buyer arrival. Examples of omitted predictors include the floor on which apartments are located and indicators for renovation needs. Other omitted factors that may help buyers in their search for housing units, aside from list prices, include those related to marketing efforts. Moreover, buyer characteristics should make it possible to approach idiosyncratic match qualities. Also, a different approach to this problem is to use buyer search history to untangle how important the list price filters are when they search online. These filters may be relevant in combination with other filters, which is an angle to the problem that leans towards segmented search (Piazzesi et al., 2020). High definition buyer search data could be the missing link that future research should focus on acquiring and utilizing.

References

Albrecht, J., Gautier, P. A., & Vroman, S. (2016). Directed search in the housing market. *Review of Economic Dynamics*, *19*, 218–231.

Andersen, S., Badarinza, C., Liu, L., Marx, J., & Ramadorai, T. (2022). Reference Dependence in the Housing Market. *American Economic Review*, *112*(10), 3398–3440.

Anenberg, E., & Bayer, P. (2020). Endogenous sources of volatility in housing markets: The joint buyer–seller problem. *International Economic Review*, *61*(3), 1195–1228.

Anglin, P. M., Rutherford, R., & Springer, T. M. (2003). The trade-off between the selling price of residential properties and time-on-the-market: The impact of price setting. *The Journal of Real Estate Finance and Economics*, *26*, 95–111.

Anundsen, A. K., Nenov, P., Larsen, E. R., & Sommervoll, D. E. (2022). *Pricing and incentives in the housing market* (Housing Lab Working Paper No. 3). Oslo Metropolitan University.

Banfi, S., & Villena-Roldan, B. (2019). Do high-wage jobs attract more applicants? Directed search evidence from the online labor market. *Journal of Labor Economics*, *37*(3), 715–746.

Barnwell, J.-L., & Fournel, J.-F. (2022, April). *Seller's (Mis)Fortune in the Housing Market: Directed Search in Online Real Estate Platforms* (Working paper). Retrieved May 8, 2024, from https://www.jeanfrancoisfournel.com/uploads/directed_search.pdf

Carrillo, P. E. (2012). An empirical stationary equilibrium search model of the housing market. *International Economic Review*, *53*(1), 203–234.

Chen, T., & Guestrin, C. (2016). Xgboost: A scalable tree boosting system. *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, 785–794.

Chen, Y., & Rosenthal, R. W. (1996). Asking prices as commitment devices. *International Economic Review*, 129–155.

Copernicus Climate Change Service. (2024a). *ERA5-Land hourly data from 1950 to present. Copernicus Climate Change Service (C3S) Climate Data Store (CDS)*. Retrieved February 26, 2024, from https://cds.climate.copernicus.eu/cdsapp#!/dataset/reanalysis-era5-land?tab=overview

Copernicus Climate Change Service. (2024b). *Nordic gridded temperature and precipitation data from 1971 to present derived from in-situ observations. Copernicus Climate Change Service (C3S) Climate Data Store (CDS)*. Retrieved February 23, 2024, from https://cds.climate.copernicus.eu/cdsapp#!/dataset/insitu-gridded-observations-nordic?tab=overview

Eiendomsmeglingsloven. (2007). Lov om eiendomsmegling (LOV-2007-06-29-73). https://lovdata.no/dokument/NL/lov/2007-06-29-73

Engelhardt, B., & Rupert, P. (2017). Competitive versus random search with bargaining: An empirical comparison. *Labour Economics*, *48*, 183–197.

Faberman, R. J., & Menzio, G. (2018). Evidence on the Relationship between Recruiting and the Starting Wage. *Labour Economics*, *50*, 67–79.

Genesove, D., & Han, L. (2012). Search and matching in the housing market. *Journal of Urban Economics*, *72*(1), 31–45.

Gourley, P. (2021). Curb appeal: How temporary weather patterns affect house prices. *The Annals of Regional Science*, *67*(1), 107–129.

Guren, A. M. (2018). House price momentum and strategic complementarity. *Journal of Political Economy*, *126*(3), 1172–1218.

Han, L., & Strange, W. C. (2016). What is the role of the asking price for a house? *Journal of Urban Economics*, *93*, 115–130.

Haurin, D. R., Haurin, J. L., Nadauld, T., & Sanders, A. (2010). List prices, sale prices and marketing time: An application to us housing markets. *Real Estate Economics*, *38*(4), 659–685.

Head, A., Lloyd-Ellis, H., & Stacey, D. (2018). *Inequality, Frictional Assignment and Home-ownership* (Working Paper No. 1396). Queen's Economics Department. https://doi.org/10.22004/ag.econ.274722

Head, A., Lloyd-Ellis, H., & Sun, H. (2014). Search, liquidity, and the dynamics of house prices and construction. *American Economic Review*, *104*(4), 1172–1210.

Head, A., Sun, H., & Zhou, C. (2016). *Default, Mortgage Standards and Housing Liquidity* (Working Paper No. 1359). Queen's Economics Department. https://doi.org/10.22004/ag.econ.274685

Hodson, T. O. (2022). Root mean square error (rmse) or mean absolute error (mae): When to use them or not. *Geoscientific Model Development Discussions*, *2022*, 1–10.

Jiang, E. X., Kotova, N., & Zhang, A. L. (2024, March). *Liquidity in Residential Real Estate Markets* (Working paper).

Lantz, B. (2023). *Machine Learning with R* (4th ed.). Packt Publishing Ltd.

Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. *Advances in neural information processing systems*, *30*.

Marinescu, I., & Wolthoff, R. (2020). Opening the black box of the matching function: The power of words. *Journal of Labor Economics*, *38*(2), 535–568.

Moen, E. R. (1997). Competitive search equilibrium. *Journal of political Economy*, *105*(2), 385–411.

Moen, E. R., Nenov, P. T., & Sniekers, F. (2021). Buying first or selling first in housing markets. *Journal of the European Economic Association*, *19*(1), 38–81.

Muñoz Sabater, J. (2019). *ERA5-Land hourly data from 1950 to present. Copernicus Climate Change Service (C3S) Climate Data Store (CDS)*. https://doi.org/10.24381/cds.e2161bac

Ngai, L. R., & Sheedy, K. D. (2020). The decision to move house and aggregate housing-market dynamics. *Journal of the European Economic Association*, *18*(5), 2487–2531.

Novy-Marx, R. (2009). Hot and cold markets. *Real Estate Economics*, *37*(1), 1–22.

Piazzesi, M., Schneider, M., & Stroebel, J. (2020). Segmented housing search. *American Economic Review*, *110*(3), 720–759.

Rekkas, M., Wright, R., & Zhu, Y. (2022, July). *How Well Does Search Theory Explain Housing Prices?* (Working paper). Available at SSRN 3706329.

Steurer, M., Hill, R. J., & Pfeifer, N. (2021). Metrics for evaluating the performance of machine learning based automated valuation models. *Journal of Property Research*, *38*(2), 99–129.

Tveito, O., Førland, E., Heino, R., Hanssen-Bauer, I., Alexandersson, H., Dahlström, B., Drebs, A., Kern-Hansen, C., Jónsson, T., E., V.-L., & Westman, Y. (2000). *Nordic Temperature Maps* (tech. rep. No. 9/00 KLIMA). Norwegian Meteorological Institute.

Tveito, O. E., Bjørdal, I., Skjelvåg, A. O., & Aune, B. (2005). A GIS-based agro-ecological decision system based on gridded climatology. *Meteorological Applications*, *12*(1), 57–68.

Wheaton, W. C. (1990). Vacancy, search, and prices in a housing market matching model. *Journal of political Economy*, *98*(6), 1270–1292.

Wright, R., Kircher, P., Julien, B., & Guerrieri, V. (2021). Directed search and competitive search equilibrium: A guided tour. *Journal of Economic Literature*, *59*(1), 90–148.

Yavas, A., & Yang, S. (1995). The strategic role of listing price in marketing real estate: Theory and evidence. *Real estate economics*, *23*(3), 347–368.

A Data description

A.1 Initial data cleaning

The data cleaning goes as follows. I restrict the data to on-market transactions in Oslo in the period September 2018 to May 2023. I keep well-behaved *transactions*, all of which must satisfy the following conditions. In some sales, a bid smaller than the maximum bid is accepted by the realtor. This can happen if the bid is placed with a condition not acceptable to the seller; if someone places a bid before getting a notice that the current bid is accepted; if the highest bid was accepted but the buyer did not have their financing in order; or if bids received days or weeks before have been rejected by the seller but the seller now has changed their mind and will not accept such a bid. Due to such potential cases, which can inflate the number of bidders, I keep units with only one accepted bid which has to be the last bid received. There are also instances where the first received bid is at least 10 percent above the list price. This can happen if it proceeds a negotiation not put on record as bids, or it may be a preemptive sale. These cases are all dropped. To keep well-behaved *auctions* without inflating the number of bidders, the auctions must end within 8 days after the first bid is received, noting that an auction ends when the seller accepts a bid.

Furthermore, to make the units have reasonable and comparable exposure on the market without any hard-to-sell units, time-on-market (TOM) is restricted at the most 31 days (one month). The TOM restriction is not particularly strict since the 95th percentile in the sample is 25 days. To make sense of the ad clicks I keep only units with one internet ad in total. And to be able to model the showings, I require there to be at least one showing.

The first bid relative to the list price is used in modeling the number of unique bidders. Some of the first bids are very low and are either typos or people wanting to register as a bidder early, being that all bidders must identify themselves with their national identification number before placing bids. I restrict the first bid to be at least 70 percent of the list price, noting that the 1st percentile is at 72 percent. Finally, the sample is trimmed on list price, ad clicks per day listed on the market, showing participants, interested people, and number of bidders.[19] Note that all cleaning steps are provided in Table A.1.

In order to add more *signal* to the models I add distances to amenities and neighborhood characteristics. Coordinates of unit locations are obtained from geocoded housing unit street addresses. These are used to calculate travel distances between the coordinates

[19] Showing participants, interested people, and number of bidders are trimmed on the 97.5 percentiles, while list price and ad clicks are trimmed on the 2.5 and 97.5 percentiles.

and amenities. I find walking distance to parks, primary schools, kindergartens, public housing addresses, tram stops, subway stations, and points of interest along the coast line. For grocery stores and driving entrances to the main highway that goes through Oslo (Ring 3 and E6), I calculate travel distance by car. I find the distances, as the crow flies, to the forest areas Nordmarka and Lillomarka, the coastal line, and areas with high degree of noise pollution from traffic, trains, subway, and trams. All these distances capture closeness to amenities that may be important for arrival, such as distance to tram stops which could be important especially when very close to a tram stop, making the housing unit either less attractive due to noise and tram passengers occupying the streets or more attractive due to closeness to public transport. See Table A.2 in the Appendix for sources of the different amenity locations, and A.3 for summary statistics of distances and elevation.

The neighborhood characteristics are collected from the City of Oslo on the *basic statistical unit* (BSU) level[20], which is a small geographical area. For each BSU, I calculate annual population density, share of population with different country backgrounds (first generation only), total housing units as a fraction of population, number of each housing unit type as a fraction of population, and the number of each housing unit type as a fraction of the total number of housing units.[21] All these variables are meant to capture the type of area of which the unit is located in. Adding to this, characterizing the surrounding BSUs, meant to represent a larger geographical area not constrained by zip code boundaries, could be as important for the popularity of housing listings. I calculated the minimum, mean, maximum, and range (max.-min.) of the neighboring BSUs for these BSU level variables, in which a neighbor j of BSU i is identified if j shares boarder with i. Including moments of surrounding BSUs should capture spatial correlations of the variables.

In addition to these characteristics, I calculate the annual frequency of which housing units are transacted on the market in each BSU, and the same types of moments for neighboring BSUs as for the BSU characteristics are also calculated. The BSU level data for housing stock and demographics are gathered from the City of Oslo and are beginning of year measures, while the market transactions are end of year measures provided by Boligmappa AS, a private company. When there is high building activity in a BSU, this may inflate the share of sold units, but the different moments for neighboring BSUs add information about these issues which allow the models to account for them.[22]

[20] In Norwegian called *grunnkrets*.

[21] These measures could all be captured by BSU fixed effects, but there are few observations in each BSU. The number of housing units in the BSU are typically between 500 to 1,500.

[22] The alternative of using the one year lag of the data from the City of Oslo will lead to missing values in the data, thus, I chose not to do this.

Most of the variables described can be characterized as hedonic variables: they characterize the housing units by their observable qualities, at the unit level and the location level. However, to capture how opportunity costs and mood affect arrival, I include weather data. Two different dataset are used, both retrieved from the Copernicus data store (Copernicus Climate Change Service, 2024a, 2024b). First, consolidated in-situ data of daily temperatures (minimum, mean, maximum) and precipitation (accumulated) at a one by one kilometer resolution (O. Tveito et al., 2000; O. E. Tveito et al., 2005). And second, a short forecast dataset of snowfall, snow depth, and wind on a roughly 9 by 9 kilometer resolution (Muñoz Sabater, 2019). These data are collected for the sample period, then summarized at the zip code level in which the housing units are located. They are included both as raw variables and as means and standard deviations in certain time frames relevant to the arrival process. The time frames are the three first days after listing, the three days before sale, the days between the final showing and the sale, and as moving averages of the last seven days before listing, final showing and sale.

The raw weather data are summarized by each of the most relevant dates, that being listing, final showing and sale, respectively, in Tables A.10, A.11, and A.12. Note that temperatures and snow variables, and partly precipitation, will capture seasonality. And because seasons vary in length and intensity over time, they will capture other seasonal effects than the usual time fixed effects which are *calendar-constrained*. For instance, while the winter is conventionally defined by calendar months (December to February), some winters can be regarded as starting earlier than December due to early arrival of low temperatures and snow, potentially persisting into March. In short, the weather does not care about the calendar.

Summary statistics of the main variables are provided in Table 1. The four main outcome variables are the *ad clicks per day* in thousands, *showing participants*, *interested people*, and *bidders*. The ad clicks are the daily averages, taking the total number of clicks and dividing by TOM, accounting for exposure time. Therefore, units with longer TOM should by construction have fewer clicks per day since the internet ads appear on top of the internet marketplace Finn.no when listed, but new listings push the ads down from their initial top positions making them less visible. This is illustrated in Table A.13, consisting of daily ad click moments in the first week. These time series are only available for a smaller subsample, which is why they are not used in this paper. The alternative to dividing ad clicks by TOM is to use the total ad clicks. When a highly popular unit is sold within a short amount of time and accumulates many clicks, this unit would appear similar to a unit that has been listed for a long time and accumulated the same amount of clicks. By dividing by TOM, how long it takes to sell the unit will be accounted for directly into the outcome variable, which is why this alternative was chosen.[23] Another alternative could be to divide the total ad clicks by time between the listing and the last showing, but this will distort the ad clicks differently based on how long between the last showing and the sales, because the available ad clicks are accumulations over the total listing period.

Showing participants include the number of people attending the showings. These are those showing up at the showings, and the same people may attend multiple showings of the same housing unit. For this second stage of arrival, how many unique people actually end up at the showings also differ from how many people are interested in the units. The variable for *interested people* includes the number of unique people registered by the realtor as being highly interested in the housing unit: they have signed up to being informed about incoming bids. These are either people who have been in direct contact with the realtor and have expressed their interest; they are people pre-registered to the showings; or they signed up as being interested at the showings. In the lockdown periods during the Covid-19 pandemic, realtors made it mandatory to pre-register for showings, while this has been a voluntary practice ever since. The interested people are informed about incoming bids and can be considered as *potential* auction participants.

The *bidders* are the unique number of bidders in each auction. Although the data does not contain identifiers for the interested people, just the total, it is not possible to check whether there are outsiders not registered as interested that participate in the auctions. However, it is highly likely that those placing bids in the auctions are all registered as interested. It is therefore reasonable to assume that bidders are among those registered, and the registered are among those clicking on the internet ad. This follows the sequence stipulated in Figure 1.

A.2 Data preparation for machine learning

The sample consists of both numeric and categorical variables. The numeric variables list price, list price per square meter (of living area), first bid divided by list price, some attributes, and distances are included with the power of one to three. Categorical variables are dummy encoded, and to deal with high cardinality in some of these variables, categories with less than 50 cases are combined into a new lumped category.[24] For each outcome, the sample is split into a 75 percent training sample and a 25 percent test sam-

[23]The TOM could have been used for modeling the total ad clicks, but I decided against this due to arrival sequence as presented in Figure 1, in which the TOM is determined by the accepted offer in the last stage.

[24]I do not do this lumping for the categorical variables year-by-month, 3-digit zip codes, and number of bedrooms. Here, I simply drop the dummies with less than 50 cases after dummy-encoding the categorical variables.

ple using *stratified sampling* of the respective outcome, so that the outcome have roughly the same distributional properties in the two samples.

Numeric variables of weather that are included as different moving averages of sometimes overlapping time periods, and BSU-level variables that are included for the respective BSU and the as different moments of the neighboring BSUs, will have cases of high correlation. This is not optimal for regression-based methods because these variables capture the same information resulting in high collinearity. However, this does not matter for tree-based methods. Still, dimensionality reduction could help performance and speed. Therefore, for the sets of weather variables and BSU variables in the training sample, I separately conduct a principal component analysis (PCA) on these sets to reduce the number of variables and replace them by uncorrelated principal components (PC). I require the PCs to capture 95 percent of the variation of the input data. The PCA treatment from the train sample is applied on the test sample.

B Summary Statistics

Table A.1: Steps taken in cleaning the sample

Cleaning step	Auctions	One bidder auctions	Multi bidder auctions
Initial	17,860	5,154	12,706
Keep market sales	17,739	5,118	12,621
Drop units sold before 2018	17,738	5,118	12,620
Keep second-hand market realtor offices	17,726	5,108	12,618
Drop units with more than one accepted bid	17,568	5,089	12,479
Drop units with bids missing their time-stamp	17,004	4,858	12,146
Drop units that received secret bids	16,898	4,823	12,075
Drop units with list or sales price being zero	16,891	4,817	12,074
Drop bids lower than one million nok	16,875	4,830	12,045
Drop counter bids	16,859	4,821	12,038
Drop units with the last bid not being the winning bid	14,925	3,888	11,037
Drop the first bid if it is 10% above the list price	14,886	3,851	11,035
Keep units sold within 8 days from the first bid is received	13,160	3,707	9,453
Drop units with a time-on-market higher than 31 days	10,694	2,635	8,059
Drop units with more than one internet ad	10,666	2,628	8,038
Drop units with zero showings	10,618	2,596	8,022
Drop units with a first bid less than 70% of the list price	10,533	2,590	7,943
Trimming the data	8,307	2,005	6,302
Drop units with the last showing not between the listing and the sale	8,273	1,975	6,298
Drop geographical outliers	8,271	1,975	6,296

Notes: The table presents the steps taken in cleaning the sample. The steps are taken to keep well-behaved transactions. Total ad clicks, divided by time-on-market (ad clicks per day), and list price are trimmed on the 2.5 and 97.5 percentiles. Showing participants, interested people, and the number of bidders are trimmed on the 97.5 percentiles. Geographical outliers indicate that these are units outside the city but within the municipality. These are located in the forest Nordmarka and are dropped.

Table A.2: Sources of added data

Data	Primary source	Secondary source	Date
Kindergartens	The Norwegian Directorate for Education and Training	Barnehagefakta	09.10.23
Grocery Stores	Kassalapp API	Tadata API	12.10.23
Primary Schools	The Norwegian Directorate for Education and Training	Geonorge	03.11.23
Public Housing	City of Oslo	City of Oslo	11.10.23
Parks	Norwegian Environment Agency	Geonorge	03.11.23
Coastal Line (points of interest)	Manually	Google Maps	13.10.23
Highway Entrances (Ring 3 and E6)	Manually	Google Maps	13.10.23
Subway stations	Ruter	Wikipedia (scraped)	10.10.23
Tram Stops	Ruter	Wikipedia (scraped)	10.10.23
Nordmarka/Lillomarka	Norwegian Mapping Authority	Geonorge	31.10.22
Coastal line	Norwegian Mapping Authority	Geonorge	31.10.22
L_{den} areas (noise)	Norwegian Environment Agency	Norwegian Environment Agency	01.03.24
Danger zones, flood and quick clay	Norwegian Water Resources and Energy Directorate	Geonorge	02.03.24
Elevation	The Norwegian Mapping Authority	Geonorge	02.10.23
Transactions, annual time series	Boligmappa AS	Boligmappa AS	07.12.23
Demographics, BSU level	Statistics Norway	City of Oslo	28.11.23
Housing stock, BSU level	Statistics Norway	City of Oslo	28.11.23
Height of buildings, # of floors	The Norwegian Mapping Authority	Geonorge	02.03.24

Notes: The table presents the data sources for amenity locations in Oslo. Primary sources are the original origins of the data, while secondary sources refer to were I collected the data from. The source Geonorge is a service provided by The Norwegian Mapping Authority, the Norwegian public mapping agency, which collects map data from different public sources. The last column *Date* reports the date in the day-month-year format.

Table A.3: Summary of shortest distances, height of close buildings, and elevation

Feature	Mean	SD	Median	Min	Max	1st quartile	3rd quartile	Type
Kindergartens	0.277	0.165	0.252	0.001	2.256	0.160	0.360	Walking
Grocery Stores	0.468	0.397	0.357	0.000	4.030	0.189	0.626	Driving
Primary Schools	0.526	0.276	0.490	0.029	2.592	0.324	0.668	Walking
Public Housing	0.658	0.458	0.540	0.000	3.276	0.299	0.931	Walking
Parks	0.394	0.259	0.334	0.000	2.146	0.214	0.502	Walking
Coastal Line (points of interest)	5.495	3.455	4.452	0.040	14.959	2.627	7.529	Walking
Highway Entrances (Ring 3 and E6)	3.205	2.052	2.701	0.102	11.072	1.874	3.855	Driving
Subway Stations	0.957	0.684	0.781	0.027	5.222	0.497	1.259	Walking
Tram Stops	2.641	2.832	1.357	0.000	10.770	0.445	4.276	Walking
Nordmarka/Lillomarka	3.656	2.613	3.220	0.012	14.224	2.116	4.365	As the crow flies
Coastal Line	3.800	2.780	2.927	0.018	12.320	1.827	4.902	As the crow flies
Closest 70 L_{den} zone (traffic)	0.155	0.152	0.111	0.000	1.642	0.042	0.219	As the crow flies
Closest 70 L_{den} zone (train)	0.366	0.367	0.269	0.000	3.523	0.111	0.511	As the crow flies
Average height of buildings (in # of floors, 50m radius)	3.993	2.102	4.000	1.000	13.000	3.000	4.700	
Elevation	0.104	0.063	0.097	0.000	0.433	0.047	0.160	

Notes: The table presents the travel distances to the nearest amenity, calculated with the Open Source Routing Machine service (https://routing.openstreetmap.de/about.html), using the R package `osrm`. All distances are provided in kilometers and in terms of walking distance, except for distance to grocery stores and highway entrances which are the driving distances. Elevation is in kilometers. The noise zones are areas with noise pollution of $L_{den} \geq 70$. The last row also reports the statistics of elevation measured in kilometers, that being the ground level elevation of the coordinates of the housing units. Note that the noise zones and danger zones for flood and quick clay landslides are included as categorical features, classifying whether the units are located in such zones. Categorical features are omitted from the table.

Table A.4: Summary of housing stock, sales rates, and demographics, BSU level, 2018

Feature	Mean	SD	Min	1st quartile	Median	3rd quartile	Max
Housing stock	662.61	275.05	110.00	471.00	628.00	817.00	1,965.00
Share sold	0.07	0.07	0.01	0.04	0.06	0.07	0.66
Share sold, weighted	0.06	0.07	0.00	0.03	0.05	0.07	0.65
Area (in 1,000 square meters)	229.44	304.00	19.92	56.66	144.71	329.86	3,824.84
Pop./Area	0.01	0.01	0.00	0.00	0.01	0.02	0.04
CB Norway	0.68	0.17	0.10	0.61	0.72	0.81	0.92
CB Western Countries	0.07	0.04	0.02	0.05	0.06	0.09	0.28
CB East Europe (EU)	0.04	0.03	0.00	0.03	0.04	0.05	0.22
CB Other	0.20	0.16	0.01	0.09	0.14	0.26	0.78
Total units/Pop.	0.52	0.13	0.24	0.42	0.52	0.61	0.98
Detached/Pop.	0.03	0.05	0.00	0.00	0.00	0.03	0.19
Semi-detached/Pop.	0.03	0.04	0.00	0.00	0.01	0.04	0.22
Row houses/Pop.	0.03	0.05	0.00	0.00	0.00	0.06	0.38
Apartments/Pop.	0.43	0.22	0.00	0.28	0.49	0.59	0.94
Detached/Total units	0.07	0.13	0.00	0.00	0.00	0.06	0.56
Semi-detached/Total units	0.06	0.10	0.00	0.00	0.02	0.09	0.55
Row houses/Total units	0.09	0.14	0.00	0.00	0.00	0.14	1.00
Apartments/Total units	0.78	0.31	0.00	0.66	0.96	1.00	1.00
# Neighbors	5.87	1.81	2.00	5.00	6.00	7.00	13.00

Notes: The table presents summary statistics on BSU level for the BSUs represented in the data, with a total of 337 BSUs. Housing stock is the total number of housing units in the BSUs, and *Share sold* is the number of market transactions as a fraction of the housing stock. The total number of transactions at the BSU level are provided by Boligmappa AS, a private company. These data do not distinguish between housing unit types, however, the housing inventory data do, so that the shares are multiplied by the apartment inventory shares resulting in the *weighted share sold*. The rest of the data are gathered from the City of Oslo, and are numbers as of January 1 annually. *CB* is an abbreviation for *country background* of the population, in which Western countries include Western Europe, the US, Canada, Australia, and New Zealand, East Europe (EU) include Eastern European countries that are members of the European Union, and Other includes all other countries. The CB statistics are shares of the total population. The population density is also calculated by the BSU area, by total housing units, and by housing unit types. The distribution of the unit types are provided by dividing on the total inventory. The last feature, $\#Neighbors$, is the number of neighbor BSUs, meaning that on average, each BSU has about 6 neighboring BSUs that they share boarders with.

Table A.5: Summary of housing stock, sales rates, and demographics, BSU level, 2019

Feature	Mean	SD	Min	1st quartile	Median	3rd quartile	Max
Housing stock	672.40	279.72	110.00	482.00	638.00	844.00	1,965.00
Share sold	0.09	0.08	0.01	0.06	0.08	0.10	0.94
Share sold, weighted	0.08	0.08	0.00	0.04	0.06	0.09	0.94
Area (in 1,000 square meters)	229.44	304.00	19.92	56.66	144.71	329.86	3,824.84
Pop./Area	0.01	0.01	0.00	0.00	0.01	0.02	0.04
CB Norway	0.68	0.17	0.10	0.61	0.72	0.81	0.91
CB Western Countries	0.07	0.03	0.02	0.04	0.06	0.09	0.19
CB East Europe (EU)	0.04	0.03	0.00	0.03	0.04	0.05	0.22
CB Other	0.21	0.16	0.01	0.09	0.15	0.26	0.76
Total units/Pop.	0.52	0.12	0.25	0.42	0.52	0.61	0.98
Detached/Pop.	0.03	0.05	0.00	0.00	0.00	0.03	0.20
Semi-detached/Pop.	0.03	0.04	0.00	0.00	0.01	0.04	0.22
Row houses/Pop.	0.04	0.05	0.00	0.00	0.00	0.06	0.39
Apartments/Pop.	0.43	0.21	0.00	0.28	0.49	0.59	0.94
Detached/Total units	0.07	0.13	0.00	0.00	0.00	0.06	0.56
Semi-detached/Total units	0.06	0.10	0.00	0.00	0.02	0.09	0.55
Row houses/Total units	0.09	0.14	0.00	0.00	0.01	0.14	1.00
Apartments/Total units	0.78	0.31	0.00	0.65	0.96	1.00	1.00
# Neighbors	5.87	1.81	2.00	5.00	6.00	7.00	13.00

Notes: The table presents summary statistics on BSU level for the BSUs represented in the data, with a total of 337 BSUs. Housing stock is the total number of housing units in the BSUs, and *Share sold* is the number of market transactions as a fraction of the housing stock. The total number of transactions at the BSU level are provided by Boligmappa AS, a private company. These data do not distinguish between housing unit types, however, the housing inventory data do, so that the shares are multiplied by the apartment inventory shares resulting in the *weighted share sold*. The rest of the data are gathered from the City of Oslo, and are numbers as of January 1 annually. *CB* is an abbreviation for *country background* of the population, in which Western countries include Western Europe, the US, Canada, Australia, and New Zealand, East Europe (EU) include Eastern European countries that are members of the European Union, and Other includes all other countries. The CB statistics are shares of the total population. The population density is also calculated by the BSU area, by total housing units, and by housing unit types. The distribution of the unit types are provided by dividing on the total inventory. The last feature, $\#Neighbors$, is the number of neighbor BSUs, meaning that on average, each BSU has about 6 neighboring BSUs that they share boarders with.

Table A.6: Summary of housing stock, sales rates, and demographics, BSU level, 2020

Feature	Mean	SD	Min	1st quartile	Median	3rd quartile	Max
Housing stock	684.40	286.95	110.00	488.00	647.00	856.00	1,981.00
Share sold	0.07	0.05	0.01	0.05	0.06	0.08	0.38
Share sold, weighted	0.06	0.05	0.00	0.03	0.05	0.08	0.38
Area (in 1,000 square meters)	229.44	304.00	19.92	56.66	144.71	329.86	3,824.84
Pop./Area	0.01	0.01	0.00	0.00	0.01	0.02	0.04
CB Norway	0.68	0.16	0.12	0.61	0.72	0.80	0.91
CB Western Countries	0.07	0.03	0.02	0.04	0.06	0.09	0.19
CB East Europe (EU)	0.04	0.03	0.00	0.03	0.04	0.05	0.20
CB Other	0.21	0.16	0.01	0.10	0.15	0.26	0.75
Total units/Pop.	0.52	0.13	0.25	0.42	0.52	0.61	1.17
Detached/Pop.	0.03	0.05	0.00	0.00	0.00	0.03	0.20
Semi-detached/Pop.	0.03	0.04	0.00	0.00	0.01	0.04	0.21
Row houses/Pop.	0.03	0.05	0.00	0.00	0.00	0.06	0.38
Apartments/Pop.	0.43	0.22	0.00	0.28	0.49	0.59	1.17
Detached/Total units	0.07	0.13	0.00	0.00	0.00	0.06	0.56
Semi-detached/Total units	0.06	0.10	0.00	0.00	0.02	0.09	0.55
Row houses/Total units	0.09	0.14	0.00	0.00	0.01	0.15	1.00
Apartments/Total units	0.78	0.31	0.00	0.65	0.96	1.00	1.00
# Neighbors	5.87	1.81	2.00	5.00	6.00	7.00	13.00

Notes: The table presents summary statistics on BSU level for the BSUs represented in the data, with a total of 337 BSUs. Housing stock is the total number of housing units in the BSUs, and *Share sold* is the number of market transactions as a fraction of the housing stock. The total number of transactions at the BSU level are provided by Boligmappa AS, a private company. These data do not distinguish between housing unit types, however, the housing inventory data do, so that the shares are multiplied by the apartment inventory shares resulting in the *weighted share sold*. The rest of the data are gathered from the City of Oslo, and are numbers as of January 1 annually. *CB* is an abbreviation for *country background* of the population, in which Western countries include Western Europe, the US, Canada, Australia, and New Zealand, East Europe (EU) include Eastern European countries that are members of the European Union, and Other includes all other countries. The CB statistics are shares of the total population. The population density is also calculated by the BSU area, by total housing units, and by housing unit types. The distribution of the unit types are provided by dividing on the total inventory. The last feature, $\#Neighbors$, is the number of neighbor BSUs, meaning that on average, each BSU has about 6 neighboring BSUs that they share boarders with.

Table A.7: Summary of housing stock, sales rates, and demographics, BSU level, 2021

Feature	Mean	SD	Min	1st quartile	Median	3rd quartile	Max
Housing stock	683.82	282.97	58.00	487.00	648.00	857.00	1,561.00
Share sold	0.07	0.04	0.02	0.05	0.06	0.08	0.42
Share sold, weighted	0.06	0.05	0.00	0.03	0.06	0.08	0.41
Area (in 1,000 square meters)	229.44	304.00	19.92	56.66	144.71	329.86	3,824.84
Pop./Area	0.01	0.01	0.00	0.00	0.01	0.02	0.04
CB Norway	0.68	0.17	0.13	0.62	0.71	0.80	0.91
CB Western Countries	0.07	0.03	0.01	0.05	0.06	0.09	0.18
CB East Europe (EU)	0.04	0.03	0.00	0.03	0.04	0.05	0.19
CB Other	0.21	0.16	0.01	0.10	0.15	0.25	0.75
Total units/Pop.	0.51	0.12	0.23	0.42	0.52	0.60	0.84
Detached/Pop.	0.03	0.05	0.00	0.00	0.00	0.03	0.19
Semi-detached/Pop.	0.03	0.04	0.00	0.00	0.01	0.04	0.22
Row houses/Pop.	0.03	0.06	0.00	0.00	0.00	0.06	0.39
Apartments/Pop.	0.42	0.21	0.00	0.28	0.48	0.59	0.84
Detached/Total units	0.07	0.12	0.00	0.00	0.00	0.06	0.55
Semi-detached/Total units	0.06	0.10	0.00	0.00	0.02	0.09	0.54
Row houses/Total units	0.09	0.15	0.00	0.00	0.01	0.15	1.00
Apartments/Total units	0.78	0.31	0.00	0.65	0.97	1.00	1.00
# Neighbors	5.87	1.81	2.00	5.00	6.00	7.00	13.00

Notes: The table presents summary statistics on BSU level for the BSUs represented in the data, with a total of 337 BSUs. Housing stock is the total number of housing units in the BSUs, and *Share sold* is the number of market transactions as a fraction of the housing stock. The total number of transactions at the BSU level are provided by Boligmappa AS, a private company. These data do not distinguish between housing unit types, however, the housing inventory data do, so that the shares are multiplied by the apartment inventory shares resulting in the *weighted share sold.* The rest of the data are gathered from the City of Oslo, and are numbers as of January 1 annually. *CB* is an abbreviation for *country background* of the population, in which Western countries include Western Europe, the US, Canada, Australia, and New Zealand, East Europe (EU) include Eastern European countries that are members of the European Union, and Other includes all other countries. The CB statistics are shares of the total population. The population density is also calculated by the BSU area, by total housing units, and by housing unit types. The distribution of the unit types are provided by dividing on the total inventory. The last feature, $\#Neighbors$, is the number of neighbor BSUs, meaning that on average, each BSU has about 6 neighboring BSUs that they share boarders with.

Table A.8: Summary of housing stock, sales rates, and demographics, BSU level, 2022

Feature	Mean	SD	Min	1st quartile	Median	3rd quartile	Max
Housing stock	689.50	288.89	58.00	490.00	653.00	871.00	1,802.00
Share sold	0.06	0.04	0.01	0.04	0.06	0.07	0.39
Share sold, weighted	0.05	0.04	0.00	0.03	0.05	0.07	0.39
Area (in 1,000 square meters)	229.44	304.00	19.92	56.66	144.71	329.86	3,824.84
Pop./Area	0.01	0.01	0.00	0.00	0.01	0.02	0.04
CB Norway	0.68	0.17	0.12	0.62	0.72	0.80	0.91
CB Western Countries	0.07	0.03	0.02	0.05	0.07	0.09	0.18
CB East Europe (EU)	0.04	0.02	0.00	0.03	0.04	0.05	0.17
CB Other	0.21	0.16	0.01	0.10	0.15	0.26	0.76
Total units/Pop.	0.51	0.12	0.22	0.42	0.52	0.60	1.01
Detached/Pop.	0.03	0.05	0.00	0.00	0.00	0.03	0.19
Semi-detached/Pop.	0.03	0.04	0.00	0.00	0.01	0.04	0.22
Row houses/Pop.	0.04	0.06	0.00	0.00	0.00	0.06	0.39
Apartments/Pop.	0.43	0.21	0.00	0.27	0.48	0.58	1.01
Detached/Total units	0.07	0.13	0.00	0.00	0.00	0.06	0.56
Semi-detached/Total units	0.06	0.10	0.00	0.00	0.02	0.09	0.54
Row houses/Total units	0.09	0.15	0.00	0.00	0.01	0.15	1.00
Apartments/Total units	0.78	0.31	0.00	0.65	0.97	1.00	1.00
# Neighbors	5.87	1.81	2.00	5.00	6.00	7.00	13.00

Notes: The table presents summary statistics on BSU level for the BSUs represented in the data, with a total of 337 BSUs. Housing stock is the total number of housing units in the BSUs, and *Share sold* is the number of market transactions as a fraction of the housing stock. The total number of transactions at the BSU level are provided by Boligmappa AS, a private company. These data do not distinguish between housing unit types, however, the housing inventory data do, so that the shares are multiplied by the apartment inventory shares resulting in the *weighted share sold.* The rest of the data are gathered from the City of Oslo, and are numbers as of January 1 annually. *CB* is an abbreviation for *country background* of the population, in which Western countries include Western Europe, the US, Canada, Australia, and New Zealand, East Europe (EU) include Eastern European countries that are members of the European Union, and Other includes all other countries. The CB statistics are shares of the total population. The population density is also calculated by the BSU area, by total housing units, and by housing unit types. The distribution of the unit types are provided by dividing on the total inventory. The last feature, $\#Neighbors$, is the number of neighbor BSUs, meaning that on average, each BSU has about 6 neighboring BSUs that they share boarders with.

Table A.9: Summary of housing stock, sales rates, and demographics, BSU level, 2023

Feature	Mean	SD	Min	1st quartile	Median	3rd quartile	Max
Housing stock	694.55	292.89	58.00	492.00	653.00	876.00	1,925.00
Share sold	0.06	0.04	0.01	0.04	0.06	0.07	0.45
Share sold, weighted	0.05	0.04	0.00	0.03	0.05	0.07	0.42
Area (in 1,000 square meters)	229.44	304.00	19.92	56.66	144.71	329.86	3,824.84
Pop./Area	0.01	0.01	0.00	0.00	0.01	0.02	0.04
CB Norway	0.67	0.17	0.12	0.61	0.71	0.79	0.91
CB Western Countries	0.07	0.03	0.02	0.05	0.07	0.09	0.19
CB East Europe (EU)	0.04	0.02	0.00	0.03	0.04	0.05	0.15
CB Other	0.22	0.16	0.01	0.10	0.16	0.27	0.77
Total units/Pop.	0.51	0.12	0.22	0.41	0.51	0.60	0.81
Detached/Pop.	0.03	0.05	0.00	0.00	0.00	0.03	0.20
Semi-detached/Pop.	0.03	0.04	0.00	0.00	0.01	0.04	0.21
Row houses/Pop.	0.03	0.06	0.00	0.00	0.00	0.06	0.39
Apartments/Pop.	0.42	0.21	0.00	0.27	0.48	0.58	0.81
Detached/Total units	0.07	0.13	0.00	0.00	0.00	0.06	0.56
Semi-detached/Total units	0.06	0.10	0.00	0.00	0.02	0.09	0.54
Row houses/Total units	0.09	0.15	0.00	0.00	0.01	0.15	1.00
Apartments/Total units	0.78	0.31	0.00	0.65	0.97	1.00	1.00
# Neighbors	5.87	1.81	2.00	5.00	6.00	7.00	13.00

Notes: The table presents summary statistics on BSU level for the BSUs represented in the data, with a total of 337 BSUs. Housing stock is the total number of housing units in the BSUs, and *Share sold* is the number of market transactions as a fraction of the housing stock. The total number of transactions at the BSU level are provided by Boligmappa AS, a private company. These data do not distinguish between housing unit types, however, the housing inventory data do, so that the shares are multiplied by the apartment inventory shares resulting in the *weighted share sold.* The rest of the data are gathered from the City of Oslo, and are numbers as of January 1 annually. *CB* is an abbreviation for *country background* of the population, in which Western countries include Western Europe, the US, Canada, Australia, and New Zealand, East Europe (EU) include Eastern European countries that are members of the European Union, and Other includes all other countries. The CB statistics are shares of the total population. The population density is also calculated by the BSU area, by total housing units, and by housing unit types. The distribution of the unit types are provided by dividing on the total inventory. The last feature, $\#Neighbors$, is the number of neighbor BSUs, meaning that on average, each BSU has about 6 neighboring BSUs that they share boarders with.

Table A.10: Summary of weather at listing dates

Feature	Mean	SD	Min	1st quartile	Median	3rd quartile	Max
Temperature (°C), daily avg.	8.31	7.22	-13.34	2.80	8.11	14.00	27.05
Temperature (°C), daily max	12.61	8.15	-12.04	6.52	12.70	19.15	32.40
Temperature (°C), daily min	4.46	6.93	-17.52	-0.49	4.11	10.09	22.09
Precipitation, daily accumulated (mm)	2.65	5.97	0.00	0.00	0.00	2.17	56.68
Snow depth (m)	0.07	0.14	0.00	0.00	0.00	0.03	0.93
Snowfall (m of water equivalent)	0.00	0.00	0.00	0.00	0.00	0.00	0.03
Wind direction (degrees)	187.17	98.83	0.01	112.06	196.01	256.60	359.98
Wind speed (m/s)	2.42	1.18	0.05	1.49	2.32	3.27	7.05

Notes: The table presents summary statistics of the weather features at the dates of listing. All features are extracted from raster data to zip codes. If a zip code overlaps with more than one raster grid cell, the average of the grid cells are used. Temperatures are in degrees Celsius (°C), precipitation is in millimeters per day (mm), snow depth is in meters (m), snowfall is in meters of water equivalent, wind direction is in degrees of which the wind comes from (north is zero, and then moving clockwise), and wind speed is in meters per second (m/s). Temperature and precipitation are based on in-situ observations, while snow depth, snowfall, and wind direction and speed are short forecasts.

Table A.11: Summary of weather at last showing dates

Feature	Mean	SD	Min	1st quartile	Median	3rd quartile	Max
Temperature (°C), daily avg.	8.31	7.04	-13.06	2.98	8.51	14.04	24.94
Temperature (°C), daily max	12.57	8.02	-9.86	6.54	12.65	19.06	31.33
Temperature (°C), daily min	4.51	6.71	-19.03	-0.16	4.34	9.90	20.70
Precipitation, daily accumulated (mm)	2.57	6.06	0.00	0.00	0.00	2.00	55.89
Snow depth (m)	0.06	0.13	0.00	0.00	0.00	0.03	0.75
Snowfall (m of water equivalent)	0.00	0.00	0.00	0.00	0.00	0.00	0.03
Wind direction (degrees)	184.27	102.33	0.00	107.70	194.90	256.89	359.93
Wind speed (m/s)	2.39	1.17	0.06	1.49	2.31	3.26	6.95

Notes: The table presents summary statistics of the weather features at the dates of listing. All features are extracted from raster data to zip codes. If a zip code overlaps with more than one raster grid cell, the average of the grid cells are used. Temperatures are in degrees Celsius (°C), precipitation is in millimeters per day (mm), snow depth is in meters (m), snowfall is in meters of water equivalent, wind direction is in degrees of which the wind comes from (north is zero, and then moving clockwise), and wind speed is in meters per second (m/s). Temperature and precipitation are based on in-situ observations, while snow depth, snowfall, and wind direction and speed are short forecasts.

Table A.12: Summary of weather at sale dates

Feature	Mean	SD	Min	1st quartile	Median	3rd quartile	Max
Temperature (°C), daily avg.	8.40	7.10	-14.77	2.88	8.45	14.07	25.16
Temperature (°C), daily max	12.64	8.14	-11.62	6.36	12.67	19.25	30.85
Temperature (°C), daily min	4.58	6.69	-17.30	-0.17	4.54	10.08	19.85
Precipitation, daily accumulated (mm)	2.60	5.96	0.00	0.00	0.00	2.14	71.33
Snow depth (m)	0.06	0.13	0.00	0.00	0.00	0.03	0.76
Snowfall (m of water equivalent)	0.00	0.00	0.00	0.00	0.00	0.00	0.02
Wind direction (degrees)	182.99	101.59	0.00	100.94	194.81	255.68	359.97
Wind speed (m/s)	2.35	1.14	0.05	1.50	2.24	3.20	7.05

Notes: The table presents summary statistics of the weather features at the dates of listing. All features are extracted from raster data to zip codes. If a zip code overlaps with more than one raster grid cell, the average of the grid cells are used. Temperatures are in degrees Celsius (°C), precipitation is in millimeters per day (mm), snow depth is in meters (m), snowfall is in meters of water equivalent, wind direction is in degrees of which the wind comes from (north is zero, and then moving clockwise), and wind speed is in meters per second (m/s). Temperature and precipitation are based on in-situ observations, while snow depth, snowfall, and wind direction and speed are short forecasts.

Table A.13: Daily ad clicks

	Frequency			Cumulative			
Day online	Min	Mean	Max	Min	Mean	Max	Obs.
1	18	1,004.88	3,583	18	1,004.88	3,583	2,231
2	111	668.41	3,653	538	1,673.29	6,258	2,231
3	0	340.09	1,450	667	2,013.38	7,702	2,231
4	0	277.72	7,206	779	2,291.10	9,006	2,231
5	0	241.63	1,986	841	2,533.08	10,289	2,229
6	0	230.00	1,644	891	2,761.46	10,586	2,221
7	0	232.69	2,036	934	2,985.03	11,502	2,173

Notes: The table presents ad clicks statistics for a subsample of the bid data in which the ad click time series are fully available. The first column *Day online* refers to the day that has passed, so that the first row represents how many ad clicks are accumulated the first day of being listed on the market.

Table A.14: Grid specifications

Method	Parameter	Values
XGBoost v1	eta	0.025, 0.05
	gamma	0.1, 0.2
	max depth	8, 10
	min child weight	2, 3
	row subsample	0.5, 0.7
	colsample by tree	0.5, 0.7
XGBoost v2	eta	0.01, 0.025
	gamma	0.05, 0.1
	max depth	6, 7
	min child weight	1, 2
	row subsample	0.5, 0.7
	colsample by tree	0.5, 0.7
XGBoost v3	eta	0.005, 0.01
	gamma	0.01, 0.05
	max depth	4, 5
	min child weight	0.5, 1
	row subsample	0.5, 0.7
	colsample by tree	0.5, 0.7

Notes: The table presents the grids used in searching for the optimal hyper-parameter combinations. There are three versions of the XG-Boost grids, *v1*, *v2* and *v3*, with the first corresponding to deeper trees and higher learning rate (eta) and the last corresponding to shallow trees and lower learning rate. The maximum number of trees to build is 3,000 (v1), 5,000 (v2), and 7,000 (v3). Early stopping is set to 50 iterations and the maximum tree thresholds are never reached.

C Results

Table B.1: Benchmark results, more metrics

	No list				List			
	MAE	RMSE	R^2	MBE	MAE	RMSE	R^2	MBE
Ad clicks per day								
OLS intercept	0.164	0.219		0.007				
OLS	0.132	0.180	0.325	0.006	0.129	0.177	0.349	0.004
Gamma	0.131	0.181	0.319	0.003	0.130	0.184	0.311	0.002
Showing participants								
OLS intercept	10.826	13.796		0.012				
OLS	8.620	11.371	0.323	0.366	8.629	11.352	0.325	0.373
NB	8.636	11.478	0.314	0.271	8.604	11.470	0.316	0.302
Interested								
OLS intercept	8.147	10.636		0.081				
OLS	6.423	8.680	0.335	0.327	6.408	8.642	0.341	0.363
Gamma	6.364	8.642	0.341	0.338	6.328	8.607	0.346	0.377
NB	6.357	8.631	0.343	0.345	6.318	8.589	0.349	0.385
Bidders								
OLS intercept	1.256	1.584		0.013				
OLS	1.221	1.562	0.034	0.029	1.195	1.535	0.064	0.031
Gamma	1.221	1.563	0.034	0.027	1.192	1.532	0.067	0.031
Poisson	1.222	1.563	0.033	0.028	1.192	1.531	0.068	0.031

Notes: The table presents results from using different regression methods for predicting the four outcome variables, with the *Ad clicks per day* being in 1,000 clicks. All methods refers to type of regression method, with *NB* referring to negative binomial regression. Note that *OLS intercept* refers to the baseline model using OLS with just an intercept term, which is equal to using the average outcome in the train sample as predictions. The R^2 is not reported for the intercept models due to no variation in the predictions. Also, the OLS intercept metrics are not repeated in the *List* group of metrics. When included, list price is included both as nominal list price and as nominal list price per square meter (size of housing unit), also taking the squared and cubic functional forms of these. All regressions include the variables living area size, age, also taking the square and cubic terms of these, ownership type, housing unit type, 3 digit zip codes fixed effects, and year-by-month fixed effects.

Table B.2: ML results with previous arrival, more metrics

Outcome	MAE	RMSE	R^2	MBE
List				
Ad clicks per day	0.114	0.164	0.454	0.020
Showing participants	7.245	9.821	0.498	0.754
Interested	4.075	5.721	0.715	0.305
Bidders	0.934	1.230	0.398	0.036
No list				
Ad clicks per day	0.117	0.167	0.433	0.019
Showing participants	7.264	9.862	0.494	0.808
Interested	4.079	5.739	0.713	0.325
Bidders	0.938	1.237	0.390	0.034

Notes: The table presents the main results of predicting the outcome variables. All models except for the ad clicks per day models use previous arrival as predictors. Ad clicks per day is in 1,000 clicks. The loss functions used in the XGBoost grid searches are the default *squared error* (all outcomes), the *Gamma* (ad clicks, interested, bidders), the *Poisson* (showing participants, interested, bidders), and the *Tweedie* (ad clicks). The gradient and the hessian of the loss function, which is part of the objective function, are used for optimizing the objective function in each iteration. The loss function for *Gamma*, *Poisson*, and *Tweedie* is the negative log-likelihood of the respective distribution.

Table B.3: ML results without previous arrival, more metrics

Outcome	MAE	RMSE	R^2	MBE
List				
Showing participants	8.340	11.158	0.351	0.895
Interested	6.081	8.490	0.382	1.242
Bidders	1.162	1.547	0.078	0.285
No list				
Showing participants	8.406	11.355	0.332	1.297
Interested	6.125	8.627	0.368	1.449
Bidders	1.174	1.560	0.060	0.278

Notes: The table presents the main results of predicting the outcome variables. None of the models use previous arrival as predictors. The loss functions used in the XG-Boost grid searches are the default *squared error* (all outcomes), the *Gamma* (interested, bidders), and the *Poisson* (all outcomes). The gradient and the hessian of the loss function, which is part of the objective function, are used for optimizing the objective function in each iteration. The loss function for *Gamma* and *Poisson* is the negative log-likelihood of the respective distribution.

Figure B.1: SHAP summary for ad clicks per day (in 1,000)

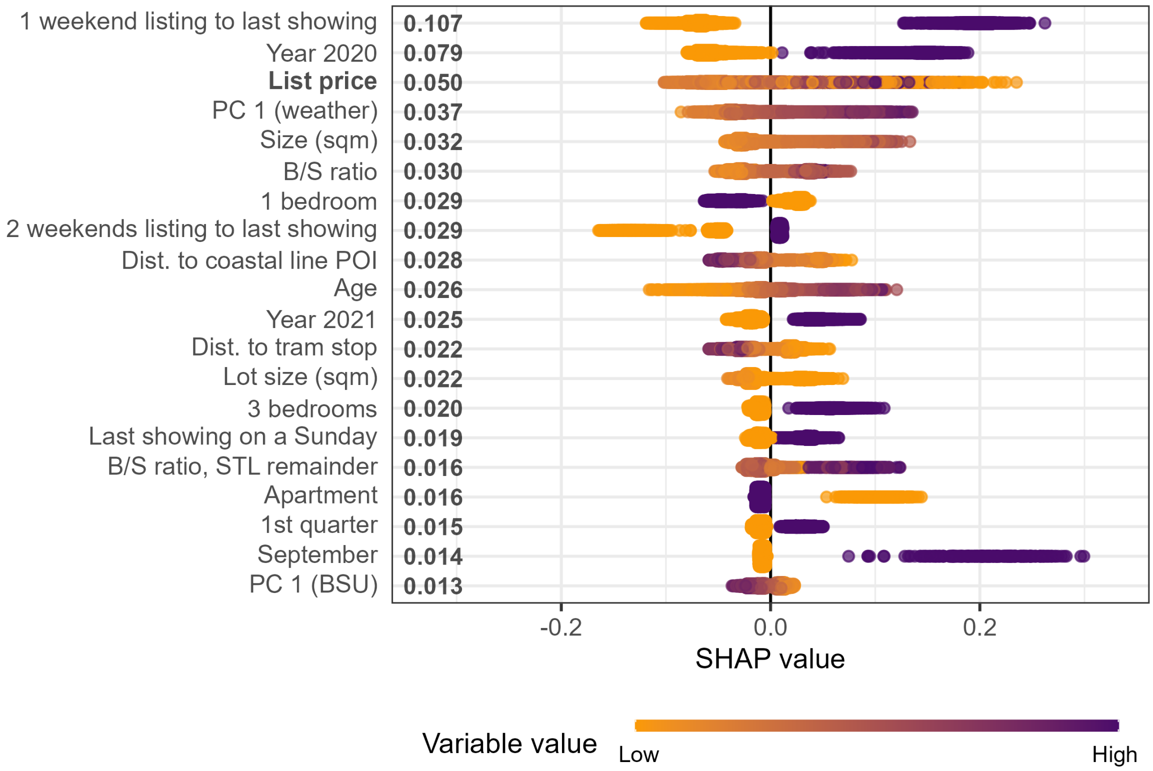

Notes: The figure presents the SHAP summary plot for ad clicks per day. The variables are ordered by the importance rank, found by taking the mean of the absolute SHAP values, provided in bold numbers. For instance, the mean absolute SHAP of the dummy for *one weekend between the date of listing and the date of the last showing* is 0.108. When this variable is 1 it has a stronger impact on ad clicks than when it is 0, as reported by the coloring of the SHAP values for this variable.

Figure B.2: SHAP summary for showing participants

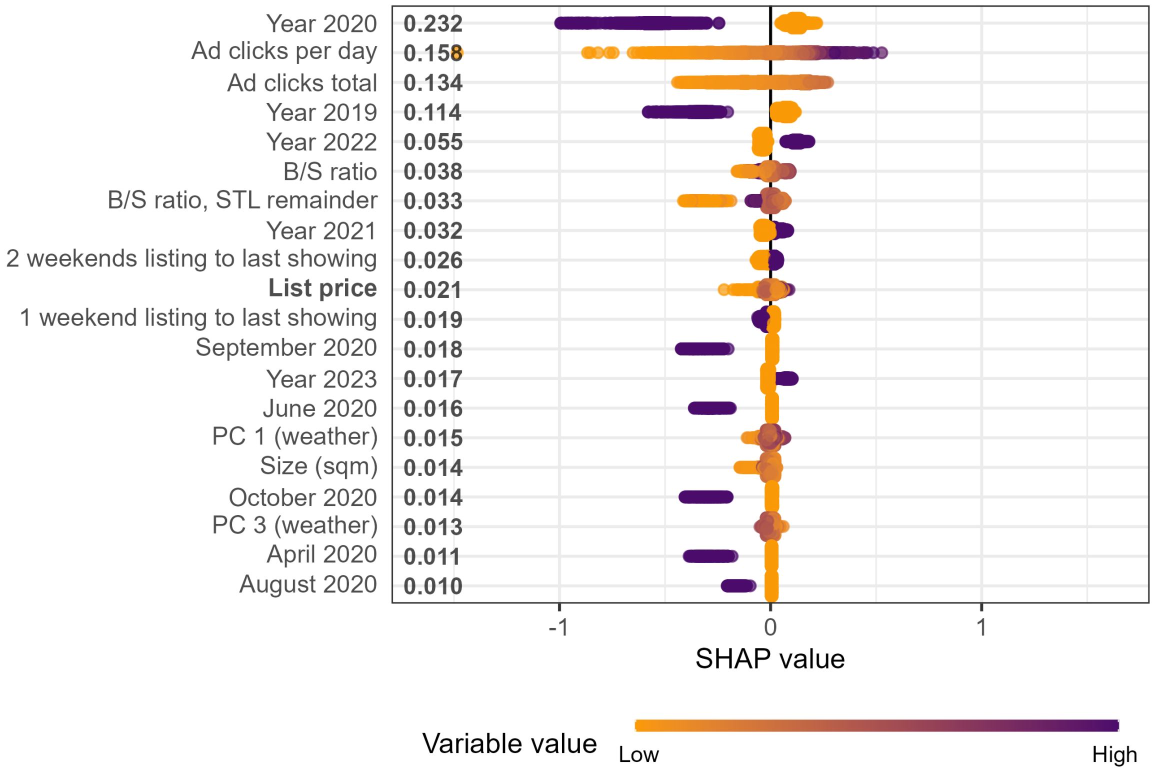

Notes: The figure presents the SHAP summary plot for showing participants. The variables are ordered by the importance rank, found by taking the mean of the absolute SHAP values, provided in bold numbers. For instance, the mean absolute SHAP of the dummy for *the sale being in 2020* is 0.235. When this variable is 1 it has a stronger impact on showing participation than when it is 0, as reported by the coloring of the SHAP values for this variable.

Figure B.3: SHAP summary for interested people

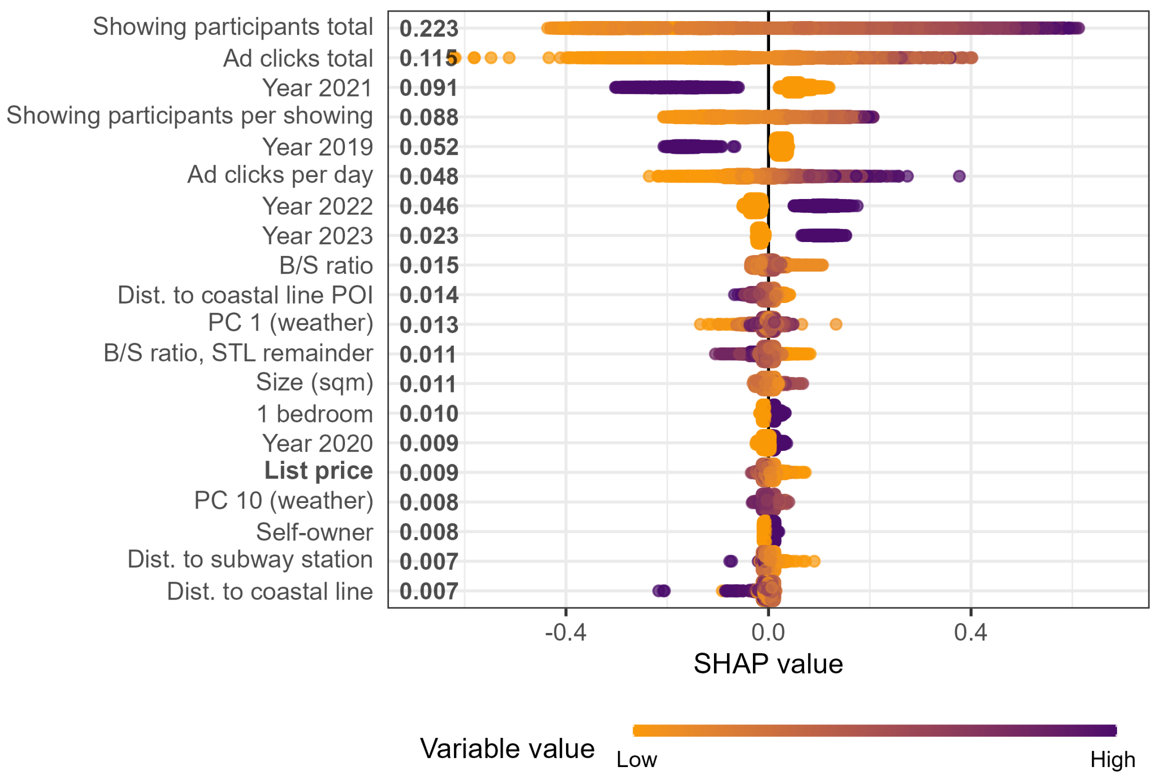

Notes: The figure presents the SHAP summary plot for interested people. The variables are ordered by the importance rank, found by taking the mean of the absolute SHAP values, provided in bold numbers. For instance, the mean absolute SHAP of the variable for *total number of showing participants* is 0.234. When this variable get closer to its extreme values it has a stronger impact on interested people, as reported by the coloring of the SHAP values for this variable.

Figure B.4: SHAP summary for bidders

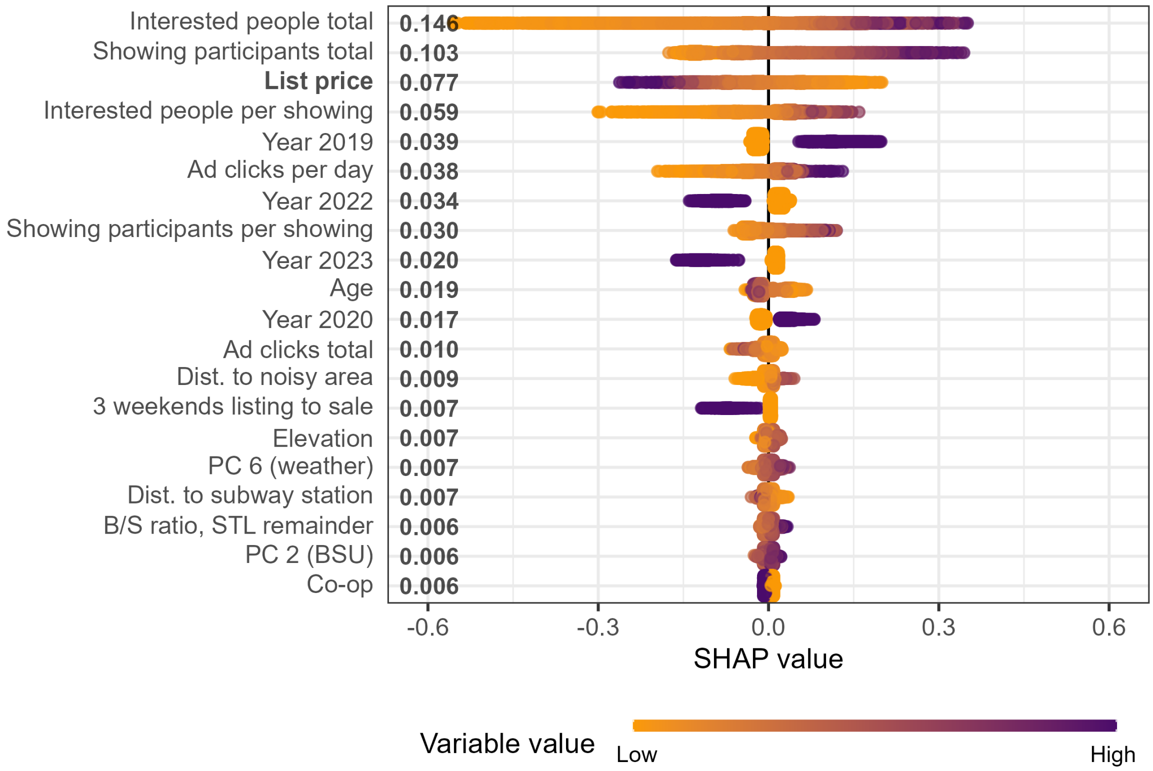

Notes: The figure presents the SHAP summary plot for bidders. The variables are ordered by the importance rank, found by taking the mean of the absolute SHAP values, provided in bold numbers. For instance, the mean absolute SHAP of the dummy for *total number of interested people* is 0.136. When this variable get closer to its extreme values it has a stronger impact on bidders, as reported by the coloring of the SHAP values for this variable.

Figure B.5: Scatter plots of outcomes against the list price

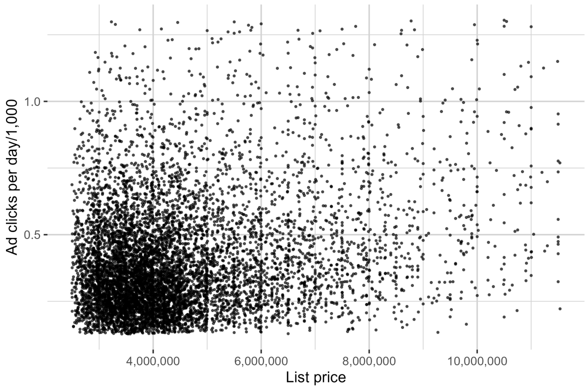

(A) Ad clicks per day/1,000

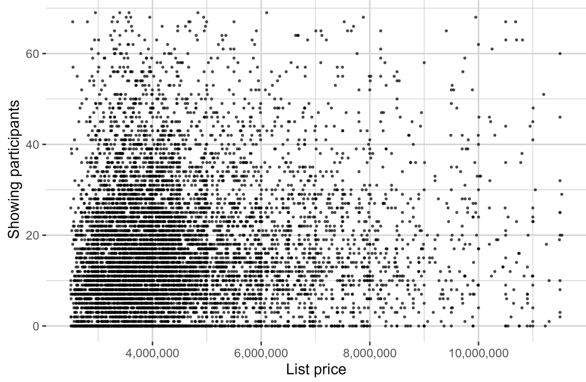

(B) Showing participants

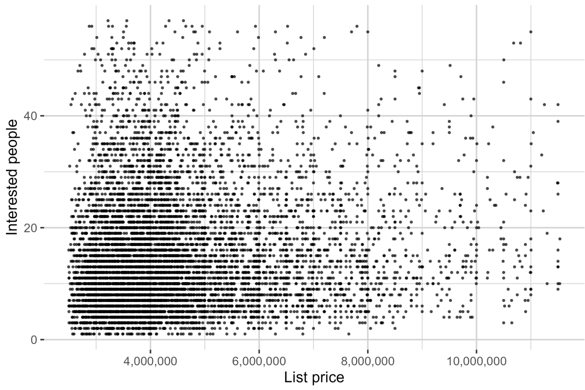

(C) Interested people

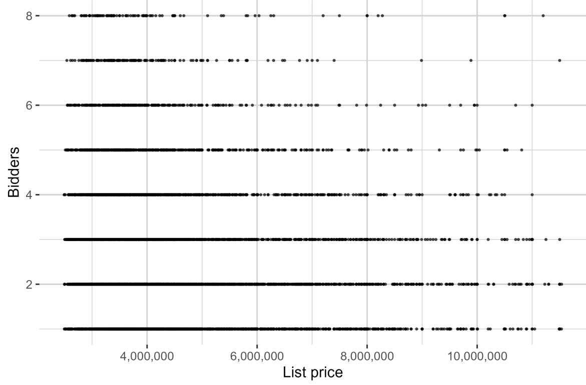

(D) Bidders

Notes: The figure presents scatter plots of the four outcome variables of interest against the list price, based on the full sample (appended train and test samples).

Figure B.6: SHAP dependence of list price for bidders, colored by arrival

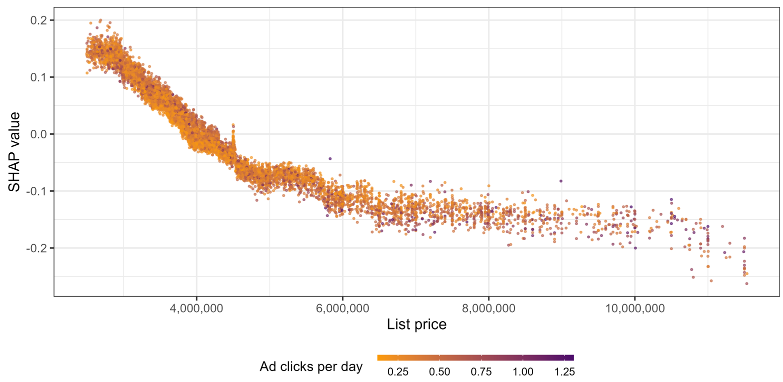

(A) Ad clicks per day/1,000

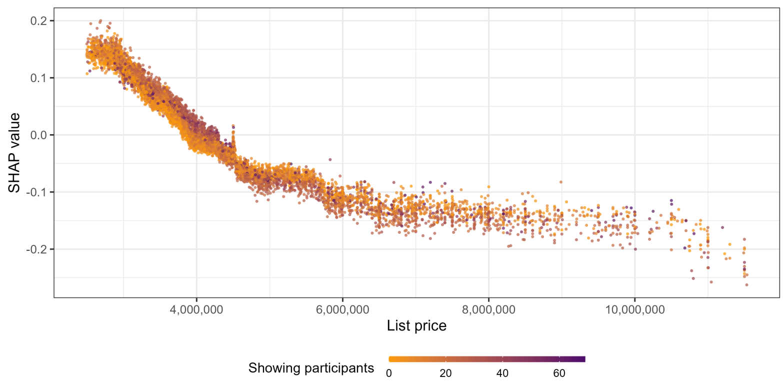

(B) Showing participants

(C) Interested people

Notes: The figure presents the SHAP dependence plot of list price for bidders. The plot is the same as Figure 3D but here I color by ad cliks per day (a), showing participants (b) and interested people (c). Positive SHAP values means that the variable affects the outcome positively, and vice versa, while the magnitude tells how strong the variable impacts the outcome.

5 Repeat bidder behavior in housing auctions

Repeat bidder behavior in housing auctions*

André K. Anundsen,[†] Andreas Benedictow,[‡] Andreas Eidspjeld Eriksen[§]
Erling Røed Larsen,[¶] and Mikkel Walbækken[‖]

November 11, 2024

Abstract

Using data on more than one million bids from almost 200,000 housing transactions, we use a unique bidder identifier to follow bidders and their bids across auctions over the period 2007–2021. We retain information on bids for each repeat bidder during their participation in different housing auctions, in addition to unit and auction-specific information. This data set allows us to estimate regression models with bidder fixed effects in order to examine whether repeat bidders change their behavior as they participate in multiple auctions. We find that they do. As the number of auctions in which they participate increase, repeat bidders bid for units that are smaller and less expensive, but bidders who eventually win their last auction tend to extend bids within their original budget range. However, since bidders' budgets remain unchanged, but the units for which they bid become smaller and less expensive, the competitiveness of their bids increase. Finally, we find that the relative bid position, as measured by the within-auction range spanned by other bids, improves when repeat bidders participate in more auctions. We detect key differences between bidders who eventually win their last auction and bidders that are not observed to win an auction.

Keywords: Auction data, bidding behavior, bid logs, housing market, repeat bidders
JEL Codes: D14; D44; D90; R21; R31

---

*We are grateful to DNB Eiendom for letting us analyze their auction data. We are grateful for comments from numerous participants in conference sessions and work-shops, including AREUEA-ASSA Annual Meeting 2023, Western Economic Association International 2022, and the National Meeting of Norwegian Resesearch Economists 2024. Thanks to Erlend Eide Bø for comments.

[†]Housing Lab, Oslo Metropolitan University.

[‡]Housing Lab, Oslo Metropolitan University and Samfunnsøkonomisk Analyse.

[§]School of Economics and Business, Norwegian University of Life Sciences and Housing Lab, Oslo Metropolitan University.

[¶]Housing Lab, Oslo Metropolitan University and Samfunnsøkonomisk Analyse.

[‖]Sopra Steria.

1 Introduction

Housing transactions are widely studied, but there is a paucity of knowledge on the behavior of bidders who experience that their bid fails and move on to bid for another house. We exploit detailed data on bidding logs and examine the behavior of bidders who bid for more than one housing unit and participate in more auctions. In Norway, a typical house sale is arranged as a digital ascending-bid auction, in which bidders compete with other bidders in real-time as they submit digital bids while the realtor informs all participants of the development. Using our access to bid logs from one of the largest realtor firms in Norway, we follow bidders across auctions over time and study the characteristics of the units for which they bid, their bids relative to the ask price of the units, their bids themselves, and their bids relative to competing bids. We aim to answer one question: Do repeat bidders change their bidding behavior as they participate in multiple auctions?

Yes, they do. We find statistically significant results that imply a change of behavior as bidders participate in more auctions. We define a repeat bidder as a buyer who is observed to extend at least one bid in at least two auctions in our bidding log data. We follow bidders across auctions until they win. Our findings indicate that repeat bidders typically seek units that are smaller and less expensive as they participate in multiple auctions. At the same time, when adjusting for time passed, they do not increase their bids but tend to stay within their original budgets. We document that compared to other competing bids, i.e., the relative position within an auction, repeat bidders make more competitive bids as they participate in more auctions. Moreover, they are observed to change behavior when we compare their bids with the ask price of the unit for which they bid and when we compare their bids to the competing bids in the auction.

Our data are sourced from bidding logs from one realtor firm, DNB Eiendom AS; a company with a market share of about 20 percent. The data highlight unique qualities of the auctions and also represent a new source of study since the data set includes all bids, not only the winning bid. Our data cover the period 2007–2021, and the trimmed data set consists of information on 1,271,741 bids, 391,305 bidders, and 195,968 transactions. Our contribution is purely empirical as we first present a few selected patterns detected in the bidding logs before we go on to examine repeat bidder behavior using regression models with bidder fixed effects.

While housing auctions in some countries may be associated with forced sales or foreclosures, the institutional arrangement of housing transactions in Norway is such that the default option is an ascending bids auction. The typical auction starts with an advertisement on the online platform Finn.no. In that advertisement, a date for the open-house is announced, and in the capital Oslo the typical advertisement is placed on the Friday nine or ten days before the open-house on the next week Sunday or Monday. The bidding starts the first Monday after the open-house and bids are extended digitally using a national, digital identifier that relies on a registry of social security numbers. However, auctions do differ in bidding frequency, bidding duration, increments of bids, and the use of bid expiration deadlines.

There are 16,480 bidders in our data set that are observed to have participated, by extending bids, in auctions of two or more units within a span of 270 days. For these bidders, we identify the highest bid in each auction within which they have participated. We first study the size of the units for which they bid. Second, we study the ask price of the units, as these ask prices relate to the annual average over ask prices in in the regions we study. Third, we study the development across auctions of the nominal bids. Fourth, we use a given bidder's highest bid and divide it by the ask price of the unit and study the development of this bid-ask ratio as bidders participate in more auctions. Fifth, we find the difference between the maximum and minimum range of bids for a given unit. Then, we compute the difference between a bidder's highest bid and the minimum bid in that auction. Subsequently, we divide the latter by the former, and use the ratio as a metric to capture a bid's relative position within the auction bid range.

We regress these five metrics on to the auction number that represents a given bidder's auction experience. For example, if a bidder participate in three auctions, we assign numbers 1, 2, and 3 to these three auctions. In our most augmented regression model, we include bidder fixed effects, number of weeks since the bidder's first bid in the first auction (in order to capture a price trend or market developments), and dummies that represent geography and bidder groups. Our main finding is that bidders seek out smaller and less expensive units, keep within their budgets, and improve the competitiveness of their bids.

Related literature

Our findings relate to the literature on the distribution of bids (Levin and Pryce (2007)), the bargaining process in the housing market (Merlo and Ortalo-Magne (2004)), auction design (Milgrom and Weber (1982), Bikhchandani and Riley (1991), Arefeva and Meng (2020) and Ettinger and Michelucci (2019)), bidding strategies (Börgers and Dustman (2005), Dodonova (2017), Hungria-Gunnelin (2018) and Sønstebø et al. (2021)), jump bids (Avery (1998), Isaac et al. (2007), Ettinger and Michelucci (2016a), Ettinger and Michelucci (2016b) and Sommervoll (2020)) and anchoring effects (Anundsen et al. (2020) and Bucchianeri and Minson (2013).

This list of studies comprises only a few contributions into the growing literature on the transaction process, bidding behavior, and auction set-up in the housing market. Since the literature on this branch of housing economics is large, we cannot here do full justice to that literature. However, let us point to a few selected studies with themes into which we contribute. An early contribution was Horowitz (1986) who constructed a bidding model in which bidders make informed choices on which auction to enter and upper limits on the bid. Although the contribution was theoretical, he also used data from Baltimore from 1978. However, the focus of attention was on the estimation of the model while our aim here is to follow bidders over time in order to investigate whether or not they change behavior. Since Horowitz (1986) many authors have examined bidding behavior, and for example Han and Strange (2014) study bidding wars in the housing market and show that they have become more frequent. This is a finding we are particularly interested in since the presence of bidding wars could lead participants to change behavior; thus the relationship to our study is clear. However, their focus of attention is on the classification of auctions while, again, ours is to inspect whether or not bidders change their bidding behavior as they gain more experience and participate in more auctions. Our finding that bids become more competitive as bidders participate in more auctions is one component that would contribute to bidding wars. Chow et al. (2015) compare auctions and negotiated sales and construct a model that has as one implication that the former tend to generate higher prices. By documenting that repeat bidders tend to change their bidding behavior in the direction of making more competitive bids, we offer empirical evidence that broadens the knowledge base of why auctions may function as a price driving set-up. This is also consistent with the findings in Arefeva (2017), who construct a search model that implies that auctions can generate higher price variance through competing bids. The dynamic created by the competition between multiple households over one unit, generates useful price signals because it shows what houses are attractive and which areas are in high demand. Genesove and Hansen (2023) show that auctions lead to prices that are useful in forecasts. We do not study the informational content generated by auctions, but we do study the mechanisms within them and the learning experience for bidders who have participated in them. So, when we inspect how repeat bidders change behavior as they participate in multiple auctions, we follow the mechanism that may lie behind the high information content in auction prices. Mateen et al. (2021) inspect 147,709 negotitations in Redfin data on how information about buyers can contribute to explain the resulting sell price by examining variables such as financing conditions, escalation clause, pre-inspection request and client letter. They find that the estimated coefficients on effects are statistically significant. From another angle Arefeva and Meng (2020) seek to understand how sellers should set a deadline for bids when auctioning their houses, a problem which illustrates the richness of the interaction between bidders and sellers. Another example of this richness can be found in Smith (2020), who constructs a model that demonstrates that there can be periods of high and low activity depending on the number of bidders a seller observes. Yet another example is Gilbukh and Goldsmith-Pinkham (2023), who use a micro data set with 8.5 million entries and show how the experience of realtors affect the sell probability. Thus, this is an active area of research, but this paper's goal of exploring change in behavior among repeat bidders appear to warrant more examination.

The structure of the article is as follows. The next section describes the institutional arrangements of the Norwegian housing market, while Section 3 presents the data set on bidding logs and the repeat bidder data set. In the subsequent section, we explain our empirical techniques. Then, we show a few graphs that captures general patterns in the data. In Section 6, we present our empirical results. In a discussion section, we perform sensitivity and robustness checks. The final section concludes the paper.

2 Institutional background

Institutional arrangements

A typical house sale process begins with a seller who contacts several real-estate agents. The realtors evaluate the market for the unit and discuss what the ask price would be. When the seller decides which realtor to hire, they find a suitable date for the unit to be put on the market and announce a date for an open-house, i.e., a public showing, in which all prospective buyers can inspect the unit. It is the seller that hires the realtor, and also the one that pays the realtor. Thus, the realtor represents the seller. However, the realtor is required by law to protect the rights of the buyer.

The advertisement, including the open-house date, is placed online with the platform Finn.no, which has a market share of at least 70 percent.[1]

An open-house is open to anyone who is interested. It provides an opportunity to inspect the home, while the realtor, and sometimes the seller, is present on location to answer

[1]https://eiendomnorge.no/housing-price-statistics/category936.html

questions from prospective buyers. These prospective buyers may also bring a professional consultant for guidance, although this is not common practice. There are strict laws governing the information duties of the seller, thus typically the advertisement includes a technical report performed by an appraiser.

The auction

In Norway, sellers typically plan to sell homes through an ascending bid auction, and only when few or no bids are obtained will the planned auction-scheme become a more stretched out negotiation process. Bids are submitted by electronic submission using digital platforms. The realtor informs the participants, as well as other people who have volunteered their name on the list of interested people, about developments in the auction.

Any bid placed is legally binding for the bidder and any accepted bid is legally binding for the seller. After a bid has been accepted, laws regulate what the seller can do to with a unit that has been sold. In essence, the legal transfer of ownership occurs on the exact hour and minute the seller accepts a bid. However, the seller may still live in the unit until the take-over date, which typically was suggested along with the bid. Since a bid may be conditional upon the take-over date, the take-over date can be seen as part of the bid. Sellers who have already bought another home typically prefer the take-over date to arrive shortly in order to avoid owning two homes and paying interest on two mortgages. Sellers who have not bought a home typically prefer the take-over date to be later in order for them to find a new home. Buyers, who are also sellers of their own unit, similarly have preferences over the take-over date. Thus, the take-over date can be a variable over which buyers and sellers negotiate. The contract signature meeting usually takes place a few weeks after the accepted bid. The registration in public registers takes place after the take-over date.

When the auction has been completed, all participants have the right to see a version of the bidding log. The log provides an overview of all the bids placed during the auction, by which stakeholder, at what time, and with what reservations. For the bidding participants and the seller the log contains actual names, in order to minimize the risk of fraudulent behavior. The names are later replaced by a scrambled ID.

Regulation of real estate brokerage and auctions

The current law on real estate entails strict requirements for the realtor, but still allows for the sales process and auction to be set up differently by different realtors, companies, and parts of the country. The law sets no requirement that homes are sold on the basis of an auction, or a specific type of auction, presence of home-seller, or home-buyer insurance, etc. Yet, nearly all ordinary sales are arranged through an ascending-bid auction. Exceptions are inheritance settlements, divorce agreements, within-family transfers etc. In recent years, regulation has been tightened. Since 2011, educational and practical requirements have been set for real estate agents and assistants, as well as for lawyers that undertake realtor assignments. As of 1 July 2010, the earliest possible acceptance deadline for a bid that is communicated via a realtor is at 12:00 PM the day after the last open-house. Since 1 January 2014, realtors can only mediate (digitally) signed bids, acceptances, or rejections.

Auction arrangement and outcomes

The Real Estate Brokerage Act gives realtors room for maneuver in designing the sales process. Therefore, it is interesting to sellers, buyers, realtors, economists, and authorities how the construction of the auction architecture affects outcomes, especially the sell price, but also the bidding process, time-on-market (TOM), and other characteristics (number of bidders, number of bids per bidder, frequencies of short expiration bids, etc.).

Realtors are required to facilitate a fair and sound settlement of the auction and adjust the pace of the bidding so that both the seller and potential buyers have a basis for acting responsibly and in line with their own interests. However, current laws and regulations do not set clear requirements for the duration of the expiration deadlines. The exception is the requirement that the realtor cannot mediate bids with a deadline earlier than 12:00 PM the first working day after the last announced open-house.[2]

In Norway, the frequency of failed sales is low. Anundsen et al. (2022) reports that Eiendomsverdi, the data analytics firm, counts sales not sold within 12 months of announcement. The rate varies by year, between three and six percent.

Conditional bids are allowed. Conditions typically involve take-over date or interior or exterior decorative elements. Legally, conditions may include many aspects.

3 Data

3.1 Bidding log data

Our data allow us to study behavioral patterns within and across housing auctions. The data set consists of Norwegian bidding log data provided by DNB Eiendom, which is the realtor arm of Norway's largest bank, DNB. The data consist of bid-by-bid logs and include detailed information of the auction, e.g., the exact time (minute resolution) when bids are placed, by whom, under what conditions, and the value of the bid. We have accessed data in two rounds, which generated one dataset from 2007 to 2018 and another that generated a dataset from 2018 to 2021. When we merge these data, we retain bidder IDs, which are unique across auctions, but an ID from the 2007-2018 data set does not transfer to the 2018-2021 data set, and vice versa.

We start out with a data set in which we have ensured that there is no missing information on the sell price, the ask price, bids, or the size of the unit transacted. Then, we remove all transactions of units that are sold more than five times over our sample period. We also truncate on the 1^{st} and 99^{th} percentiles of the sell price, ask price, and size. This truncation is done year-by-year and for nine different regions of Norway.[3]

[2] In Denmark and Sweden, home buyers in practice have more time. In Sweden, a bid is not binding, which constitutes an advantage for the highest bidder, who in the aftermath of the auction can assess the home with professional assistance and choose to withdraw the bid. In Denmark, it is possible to withdraw from the transaction against a fee that represents a percentage of the selling price.

[3] The regions consist of eight municipalities and the remaining sample: Asker, Bergen, Bærum, Fredrik-

We remove "unserious" bids – bids that are less than 60 percent of the ask price. We also identify the 99^{th} percentile of the bid-distribution, and remove the entire auction if at least one bid in the auction is greater than or equal to the 99^{th} percentile. This is done to rule out clerical errors and bids that are extreme.

Finally, we remove transactions that are recorded with a sell date that differs from the date of bid-acceptance, units for which the final bid is recorded as being accepted before it was placed, auctions for which any bid expires when or before it is placed. After these adjustments, our data set consists of 195,968 transactions and 1,271,741 bids.

We use this data set to calculate a set of transaction-specific measures: size, sell-ask spread, number of bids, number of bidders, number of bids per bidder, TOM, and bidding duration. We also calculate the spread between the first bid in each auction and the ask price. These variables are summarized in the upper part of Table 1, in which we also record the fraction of auctions that are owner-occupied and the fraction of apartments.

In the lower part of Table 1, we show the distribution of a set of auction-specific variables: expiration deadline of bids (in minutes from placing the bid), the time elapsed between the placement of subsequent bids within an auction, and the percentage change in bid increments. Finally, we present our measures of the spread between a given bidder's highest and lowest bid within the same auction, conditional on observing that the bidder makes at least two bids.

The median sell-ask spread is zero, which suggests that the ask price has predictive power for the sell price. It does, however, have a right-tailed distribution, with a positive mean sell-ask spread.

Both the mean and median number of bidders are equal to 2, but there are many auctions with both a lower and a higher number of bidders. The highest number of bidders recorded in our data set is 34. At the auction-level, the median number of bids is 5 and the highest number of bids is 84. Each bidder typically makes four bids (75th pct.) or less, but in

stad, Lillehammer, Oslo, Skedsmo, Trondheim, and the remaining sample.

multiple cases they make more bids.

90 percent of the units in our data set are sold within 54 days of the list date, and the median TOM is 11 days.

While several bidding processes are completed quickly, some have long duration. If we denote the length of time from the first bid to the accepted bid as the bidding duration, we observe that more than 75 percent of bidding processes have a duration that is less than or equal to 3 days. We observe many cases in which there are relatively short expiration deadlines, and about 75 percent of the deadlines are two hours or less. Although some bidders extend bids with long deadlines, bids are countered quickly, and the median response-time is 15 minutes.

In most auctions, the value of the first bid is well below the ask price, and the mean spread between the opening-bid and the ask price is 7 percent. Bid increments are typically relatively small, although there are cases in which increments are substantial – at the 90^{th} percentile, bids are increased by five percent relative to the previous bid.

At the median, the difference between the highest and the lowest bid of a bidder is seven percent, which suggests that bidders start their bidding well below their willingness-to-pay.

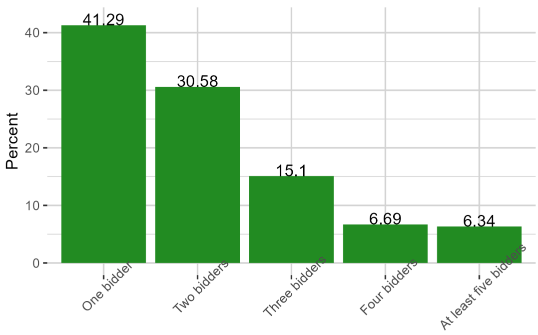

(a) Bidders

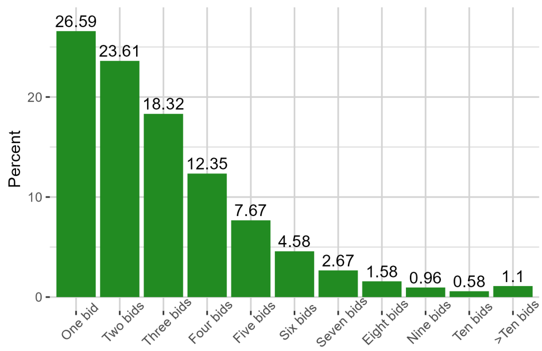

(b) Bids per bidder

Figure 1: Panel a) shows the distribution of the number of bidders in Norwegian housing auctions, 2007–2021. Panel b) shows the distribution of the number of bids placed by a given bidder in Norwegian housing auctions, 2007–2021. Calculations are based on auction-logs from the realtor-firm DNB Eiendom.

Table 1: **Summary statistics for bid-level sales data. Norway, 2007–2021**

	10^{th} pct.	25^{th} pct.	Median	Mean	75^{th} pct.	90^{th} pct.
Sale-specific						
No. bidders	1.00	1.00	2.00	2.13	3.00	4.00
No. bids	1.00	2.00	5.00	6.49	9.00	14.00
No. bids per bidder (total)	1.00	2.00	3.00	3.50	4.00	7.00
No. bids per bidder (by sale)	1.00	1.00	2.00	3.04	4.00	6.00
TOM (days)	7.00	8.00	11.00	22.87	24.00	54.00
First-bid-ask spr. (in %)	-16.32	-10.78	-6.10	-6.90	-1.91	0.25
Spread btw. highest and lowest bid (in %)	2.13	3.92	7.14	9.23	12.00	18.43
Unit-specific						
Sell-Ask spr. (in %)	-5.59	-2.61	0.00	2.12	5.56	12.71
Ask (in million NOK)	1.50	1.95	2.64	2.91	3.55	4.73
Size (in sqm.)	50.00	65.00	89.00	101.22	130.00	173.00
Owner-occupied (in %)				68.94		
Oslo (in %)				18.39		
No. bidders (total)				391,305		
No. bids (total)				1,271,741		
No. sales (total)				195,968		

Notes: The table shows summary statistics for sales-level data over the period 2007 – 2021. We have calculated the mean and median, in addition to the 10^{th}, 25^{th}, 75^{th}, and 90^{th} percentiles for each variable. All calculations are based on auction-logs from the realtor-firm DNB Eiendom. We include common debt in the computation of ask price.

3.2 Repeat bidder data

Our data set allows us to follow the same bidder across multiple auctions. Within the same auction, we have seen that a bidder's maximum bid may be considerably larger than the bidder's minimum bid. A natural extension of examining repeat-bidders within auctions is an examination of repeat-bidders across auctions. We are particularly interested in studying patterns that involve a given bidder's maximum bid in each auction the bidder participates in. To this end, we follow repeat-auction participants, which we define as bidders that have participated in at least two auctions.

We define a repeat bidder as a bidder who participate in at least two auctions within a span of 270 days. This means that they can be observed in at most two different calendar years. We differentiate between bidders we observe to win (winners) and those not winning (losers), and drop winners if we observe them placing bids after they won. In Figure 2 we show the distribution of how often we observe them – the observed experience. Most of the repeat bidder data consist of bidders observed in two auctions (88 percent) while a smaller minority are observed in three auctions (10 percent). We also show the distribution of how many bids the repeat bidders place in auctions, which closely resembles the distribution for the full sample as presented in Figure 1. All steps taken to construct the repeat bidders data are provided in Table 2. Summary statistics of this data set are provided in Table 3. Notice that we include common debt in the computation in ask price since we shall use that feature in our analysis of repeat bidders.

Table 2: Construction of the repeat bidder data set

Cleaning step	# Auctions	# Bidders	# Bidder-Auction inst.
Initial	195,968	391,305	418,010
Drop secret and counter bids	195,837	390,990	417,655
Drop auctions with 1 bid	167,326	364,275	389,144
Keep bidders observed in at least 2 auctions	36,161	20,969	45,838
Drop missing time between first and last bid observed by bidder	35,982	20,848	45,523
Drop bidders observed over a longer period than 270 days	30,649	17,617	37,988
Keep bidders observed winning at most once	29,693	17,026	36,606
Drop bidders if observed after they win	28,877	16,480	35,379
Drop bidders if observed in more than 5 auctions	28,813	16,454	35,207

Notes: The table presents the steps taken from the bid-level sales data to the repeat bidders data. The column *Bidder-Auction inst.* presents the bidder-auction instances observed, that is, the total number of bidder-by-auction cases in the sample. To keep track of the sequence at which bidders are observed, we define the bidder participation in an auction as the last time the bidder placed a bid in that auction. That means, if the bidder placed a bid in an auction, then moved on to bidding on other units while the first auction did not end with a sale, and then came back bidding again at the first unit, the first unit will be the last in this sequence.

3.3 Example bid

We present here an example of bid with example variables included. In order to ensure anonymity, we have constructed a fictitious bid that maintains the essential elements of the information in our data.

"Unit id: 10101; Assignment id: 20202; Auction First bid date: 2013-04-05; Auction Final Bid accept date: 2013-06-20; Auction Final bid expire: 2013-06-21 13:20;Present bid date:

Table 3: **Summary statistics for repeat bidder data. Norway, 2007–2021**

	10^{th} pct.	25^{th} pct.	Median	Mean	75^{th} pct.	90^{th} pct.
Sale-specific						
No. bidders	1.00	1.00	1.00	1.22	1.00	2.00
No. bids	1.00	2.00	3.00	3.52	4.00	7.00
No. bids per bidder (total)	3.00	4.00	6.00	6.36	8.00	11.00
No. bids per bidder (by sale)	1.00	1.00	2.00	2.88	4.00	6.00
TOM (days)	7.00	8.00	10.00	17.81	14.00	36.00
First-bid-ask spr. (in %)	-17.67	-11.90	-6.98	-7.82	-2.70	0.00
Spread btw. highest and lowest bid (in %)	2.14	3.95	7.22	9.38	12.28	18.92
Unit-specific						
Sell-Ask spr. (in %)	-3.23	0.00	5.00	6.72	11.93	19.09
Ask (in million NOK)	1.60	2.04	2.64	2.85	3.40	4.34
Size (in sqm.)	43.00	56.00	74.00	88.74	113.00	155.00
Owner-occupied (in %)				62.60		
Oslo (in %)				29.04		
No. bidders (total)				16,454		
No. bids (total)				101,280		
No. sales (total)				28,813		

Notes: The table shows summary statistics for repeat bidder data over the period 2007 – 2021. We have calculated the mean and median, in addition to the 10^{th}, 25^{th}, 75^{th}, and 90^{th} percentiles for each variable. All calculations are based on auction-logs from the realtor-firm DNB Eiendom. We include common debt in the computation of ask price.

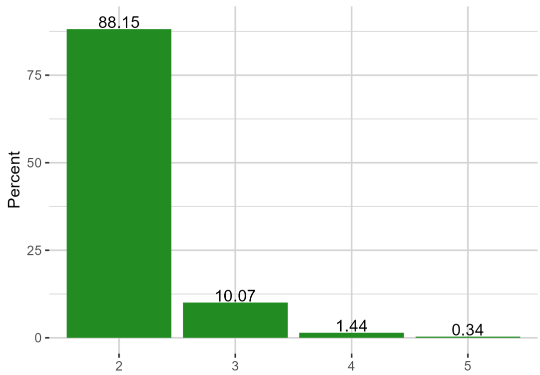

(a) Bidders per observed experience group

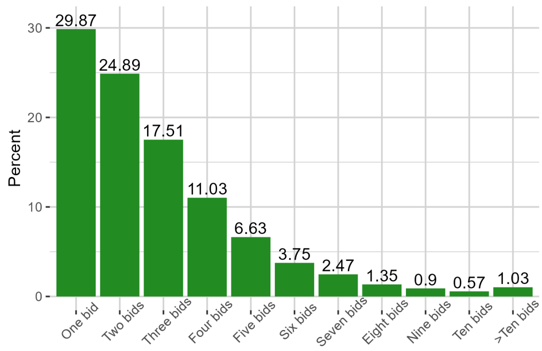

(b) Bids per bidder

Figure 2: Panel a) shows the distribution of the number of bidders per observed experience group in the repeat bidders data. Panel b) shows the distribution of the number of bids placed by a given bidder in the repeat bidders data. Calculations are based on auction-logs from the realtor-firm DNB Eiendom.

2013-04-08 14:00;Present bid expire: 2013-04-08 15:00;Present bid conditions[4]: Yes;Unit size:

[4]Exists for data 2007-2018.

93;Unit type: owner-occupied;Unit no bedroom: 3;Unit ask: 5,300,000;Present bid value: 5,200,000;Present bidder no bid: 2;Auction no bidders: 5;Auction no bids: 17;Auction min bid: 4,900,000;Auction max bid: 5,400,000."

4 Empirical techniques

Our main contribution is the regression results from testing hypotheses on how bidders behave when they participate in more than one auction. Do bidders change behavior in terms of what units they seek to buy, how their bids compare to ask price, and how their bids compare to competitors? It is an empirical question whether bidders do change behavior as they participate in more auctions since both a finding of no change in behavior and a finding of changed behavior can be explained by theory. No change in behavior would be consistent with a theory of forward-looking, fully informed individuals who have acquired the required knowledge base before the first auction. A change in behavior would be consistent with a theory in which learning is a key component.

This article studies the choice of units, a bidder's bids compared to ask price, a bidder's nominal bid, and a bidder's position relative to competing bidders. By a bidder's bid, we mean the highest bid a bidder places within an auction, which is different from the within-auction highest bid across bidders. We denote the within-auction highest bid "the maximum bid".

Our null hypotheses are:

1. Bidders bid on equally sized units
2. Bidders bid on equally priced units
3. Bidders place bids that are of equal values
4. Bidders place bids that have the same bid-ask spread

5. Bidders' highest bids relative to auction maximum stays the same

We test our hypotheses by inspecting the estimated coefficient of a continuous variable 'Auction number', which represents the stage at which a bidder is during the auction sequence.

$$Z_{ba} = \alpha + \beta_1 Au_{ba} + \beta_2 AuO_{ba} + \beta_3 N_{ba} + \beta_4 T_{ba} + \beta_5 O_{ba} + \beta_6 FE_b + \varepsilon_{ba}, \tag{1}$$

in which subscript ba represents a running number of bidder-auction combinations. For example, if bidder $b = A$ participates in three auctions $(a = x, y, z)$ and bidder $b = B$ participates in four auctions $(a = x, y, w, r$, there would be seven unique bidder-auction combinations. Then, ba would run from one to seven. We use five dependent variables Z, we segment into winners and losers of the last auction and we segment Oslo and non-Oslo. We introduce each determinant sequentially so that we run one new regression for each determinant introduced. This means that we estimate ten regressions for each of the five dependent variables, a total of fifty regressions. The results are presented in five tables below.

These are the five dependent variables Z we use in our four regression categories:

1. Size of the unit in auction a for which a given bidder b bids
2. The ask price of the unit in auction a for which a given bidder b bids
3. The highest bid a bidder b extends in an auction a
4. The bid a given bidder b makes compared to the ask price in auction a
5. The highest bid bidder b makes less the minimum bid in auction a divided by the difference between the maximum and minimum bids in the auction, $(bid_{ba} - min_a)/(max_a - min_a)$

These are the determinants we introduce sequentially:

1. Au. Auction number, i.e., the number of the auction in the sequence of auctions in which bidder b participates. If bidder b participates in four auctions, 'Auction number' = 3 would mean that we look at the third auction, out of four, in which bidder b participates

2. N. Short notation for "N Auctions". A vector of dummies for bidder groups, segmented by the total number of auctions in which a bidder participates. For example, if bidder b participates in two auctions, the dummy variable for "N auctions 2" is unity and "N auctions 3", "N auctions 4", and "N auctions 5 or more" are zero.

3. T. The date of the highest bid in auction number a less the date of the first bid in auction number one, measured in weeks

4. O. A dummy for Oslo

5. AuO. Interaction term $Au \times O$

6. FE. Bidder fixed effects

We use as a maintained assumption that a bidder visits units that are sold by different realtors randomly. DNB has a market share of 20 percent, thus we observe only one out of five auctions a bidder participates in. A given bidder that is observed to participate in two auctions in our dataset may have participated in eight other non-observed auctions. We assume that the participation sequencing is random between realtors. The implication is that for a participant who in total participates in ten auctions, out of which two are registered in DNB data, the sequence number of those two can be modelled as two numbers drawn randomly without replacement from ten numbers, 1 through 10.

5 Illustrative examples and patterns

Our data consist of bids, bidders, auctions, and houses; all potential units of analysis. In order to illustrate the construction of our regression analysis and to highlight patterns in the data, we present two graphs in this section. The idea is that by discussing the visual traces of bidding behavior we motivate the specification choices in our regression analyses. We do this by generating graphs along four dimensions:

- Bids relative to the ask price of the unit
- Bids relative to competing bids
- Segmentation on bidders that win the last auction and bidders that do not win any auction
- Segmentation on bidder experience; i.e., the number of auctions in which they have participated

The backdrop of this reasoning is that it is possible for unsuccessful bidders to ensure a win by looking at smaller and less expensive units while using the same buying budget, a possibility that shall be fully addressed in the regression analyses in the next section. Moreover, it is possible that more experienced bidders, i.e., bidders that we observe to participate in more auctions, behave differently compared to less experienced bidders.

Figure 3 depicts results from plotting a bidder's highest bid relative to the ask price. The plot consists of two panels, in which the left panel displays the results from analyzing the behavior of bidders that did not win an auction and the right panel displays the results from analyzing the behavior of bidders that won the last auction in which they participated.

In each panel, we present results from four groups of bidders. We classify along the dimension of their bidding experience, i.e., how many auctions in which they have been participated. For example, the two red lines are computed on data consisting of bidders that are observed in our data set to have extended bids in two auctions. The first node on the red line is the mean across bidders in their first auction. The second node on the red line is the mean across the same bidders in their second and last auction. The two blue lines are computed on data consisting of bidders that are observed to have extended bids in three auctions. The nodes represent the means, across bidders, in the first, second, and third auction, respectively.

We detect a pattern in which all eight lines have a higher end nodes than start nodes. The interpretation is that bidders tend to increase their bids compared to ask price when they participate in more auctions. The visual impression is that the increase is larger among bidders who won the last auction. However, we caution against statistical inference since these graphs do not control for bidder fixed effects nor time effects. Below, we estimate regression models that allow us to make inferences. Moreover, keep in mind that the number of observations that underlie the red lines is much larger than the number of observations that underlie the purple lines. In the regression models, this data difference will entail weighting the groups.

In Figure 4, we plot the development of how bidders position themselves in relation to other bidders within the same auction. Again, we first segment into bidders that won their last auction and bidders that did not win any auction. Then, we segment bidders into experience groups depending on the number of auctions in which the bidders are observed to participate. The red lines are computed on the basis of bidders that participated in two auctions. The purple lies are computed on the basis of bidders that participated in five or more auctions.

We identify the maximum and minimum bid in all auctions. Then we investigate how each bidder's own highest bid can be placed in relation to these two within-auction end points of the bidding range. For each bidder, we identify where that bidder's highest bid can be placed on the range from an auction's minimum bid to the auction's maximum bid. We do this by finding the ratio for each bidder b between b's highest bid in auction a less the minimum bid in the auction a on that auction's bid range, i.e., the difference between the

Figure 3: Bids relative to unit value and observed bidder experience

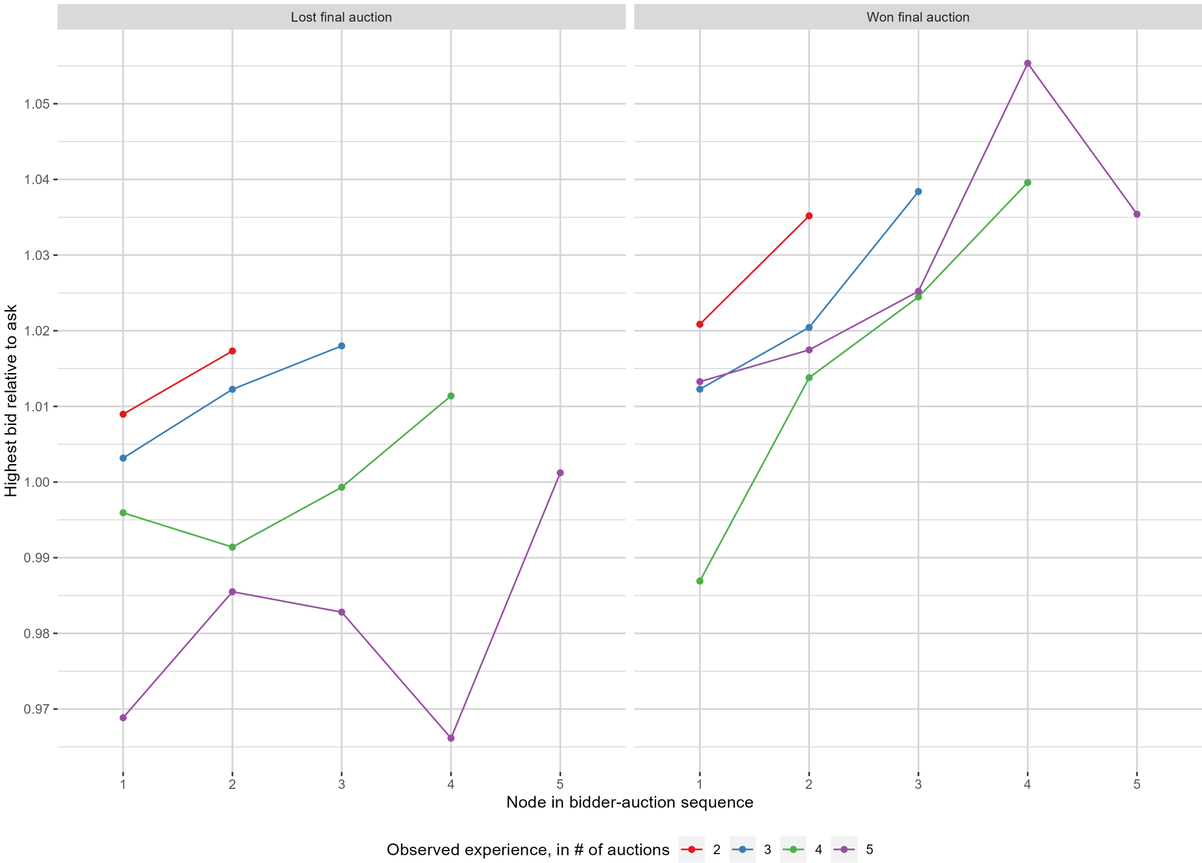

Notes: We first compute the ratio of the highest bid on the ask price for each bidder in an auction, then compute the mean of this ratio across all unique bidder-auction combinations, for each node in the auction sequence in which bidders participate. We partitioned into bidders that lost the final auction (left panel) and bidders that won the last auction (right panel).

maximum bid and the minimum bid, $(bid_{b,a} - min_a)/(max_a - min_a)$. Then, we compute the mean position, across bidders, at each node for each group.

Figure 4 shows that the slopes of the lines in the left hand-side panel are close to zero while the slopes of the lines in the right hand-side graph are clearly positive. The interpretation is that bidders that win their last auction tend to bid in a manner, before they win, that ensures that their highest bid gets closer to the winning bid. By construction, their last bid is the winning bid.

Figure 4: Bids relative to competing bids and observed bidder experience

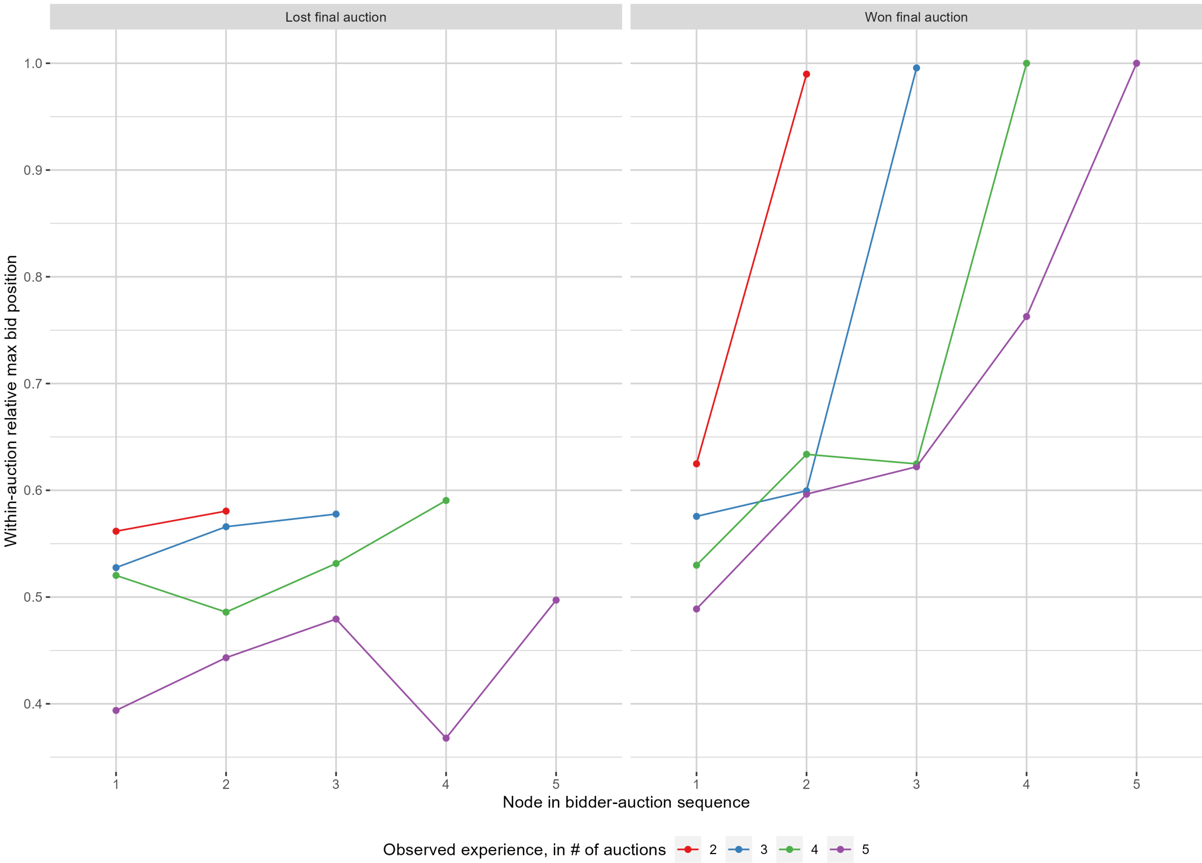

Notes: The figure presents the mean, across bidders, of the relative bid position. The relative bid position is defined as the ratio between a given bidder b's highest bid in auction a less the global minimum bid in auction a on the difference between the global maximum bid in auction a and the global minimum bid in auction a, (Highest $\text{bid}_{b,a}$ − Minimum bid_a)/ (Maximum bid_a − Minimum bid_a). We partitioned into bidders that lost the final auction (left panel) and bidders that won the last auction (right panel).

Thus, for the right-hand panel it is not surprising that there is an increase from the second-to-last node to the last node. After all, that last bid won the auction, and our partitioning implies a high last bid. It is also possible that a winner has placed a bid of the same value in several auctions and finally wins because she finds herself in a "cold" auction with little competition. If so, there would be a selection effect. However, we do observe a tendency for these bidders also to increase their relative bids in earlier auctions, i.e., they appear to change their behavior in auctions before the auction they win. This is worth investigating because it is consistent with the idea that these winners changed their behavior before they were able to win. We return to this idea in our more rigorous examination below.

Notice also that the number of observations for the purple line in the right panel is 24. Compare that number of observations to the red line, for example, which is 6,536.

6 Empirical results

6.1 A repeat bidder's choice of units with respect to size

In Table 4, we tabulate the results from ten regressions of unit size on the variable "Auction number". Unit size $S_{b,a}$ is the size of the unit for which bidder b bids in auction a. "Auction number" is a continuous variable that takes on sequential, counting numbers for each auction in which a bidder participates. For example, if bidder b participates in four auctions, this bidder would be observed in bidding in four different auctions, numbered $a = 1$, $a = 2$, $a = 3$, and $a = 4$. If another bidder c participates in three auctions, some of which could overlap bidder b, bidder c would be observed with a equal to one, two, and three. The number of auctions in which bidders participate varies across bidders. Since bidders that participate in two auctions could behave differently from bidders that participate in five auctions, e.g., due to experience, we segment on the total number of auctions in which a bidder has participated. The idea behind this segmentation is the role played by learning. Thus, there is a dummy variable that is unity for bidders that participate in three auctions and zero otherwise. Similarly, there is a dummy variable that is unity for bidders that participate in four auctions and zero otherwise. The default is two auctions. Our Model 5 includes bidder fixed effects, and in that type we cannot use the dummy vector "N auctions" due to the perfect correlation between the individual bidder dummy and the "Auction number".

We run five regressions for bidders that win the last auction (Panel A) and five regressions for bidders that lose all auctions (Panel B). The first regression is an OLS regression of unit size on auction number for all bidder-auction combinations. In the second regression, as we explained above, we include a dummy for each of the categories "Participated in three auctions", "Participated in four auctions", "Participated in five auctions" (Participated in two auctions is default). Short notation for this extension of a list of dummy variables is "N auctions". We also include a variable "Weeks from first observed" that is defined as the number of weeks between the date of the last bid in the last auction and the date of the first

Table 4: Unit size. Regressions of unit size on auction number. Partitioned into bidders that lost the last auction and bidders that won the last auction

Panel A: Winners

	(1)	(2)	(3)	(4)	(5)
Auction number	-2.250^{***}	-5.032^{***}	-3.951^{***}	-4.844^{***}	-0.766^{**}
	(0.602)	(0.737)	(0.680)	(0.906)	(0.342)
Auction number*Oslo				2.484^{***}	
				(0.953)	
N auctions		✓	✓	✓	
Weeks from first observed		✓	✓	✓	✓
Oslo			✓	✓	✓
Bidder FE					✓
Observations	12,663	12,663	12,663	12,663	12,663
Adjusted R^2	0.001	0.017	0.156	0.156	0.843

Panel B: Losers

	(1)	(2)	(3)	(4)	(5)
Auction number	-2.973^{***}	-3.940^{***}	-2.435^{***}	-3.659^{***}	-0.067
	(0.491)	(0.593)	(0.542)	(0.838)	(0.262)
Auction number*Oslo				2.503^{***}	
				(0.861)	
N auctions		✓	✓	✓	
Weeks from first observed		✓	✓	✓	✓
Oslo			✓	✓	✓
Bidder FE					✓
Observations	16,190	16,190	16,190	16,190	16,190
Adjusted R^2	0.002	0.023	0.189	0.190	0.849

Notes: The table shows results of regressing the size in square meters of the bidder-auction combination *ba* on the auction number for that combination. Panel A gives the results from the subset of bidders that win the last auction in which they participate, and panel B gives the result for those not winning the last auction. We choose to drop bidders if they are observed to bid on different unit types, e.g. first a detached house and then an apartment. If bidders are observed to win but still place bids afterwards, they are dropped from the sample. The dummy *N auctions* is a control variable for the number of auctions in which each bidder has participated, *Weeks from first observed* represents how many weeks have passed since the bidder is first observed, *Oslo* indicates whether the auction is located in Oslo, and *Bidder FE* are bidder fixed-effects. We retain bidders that have participated in five or less auctions. Heteroskedasticity-robust standard errors are given in parenthesis. Significance: * p<0.1, ** p<0.05, *** p<0.01.

bid in the first auction. This variable controls for potential short-term changes in the price level. In the third regression, we include a dummy for the auctioned unit being located in the capital, Oslo. In the fourth regression, we add an interaction term between "Auction number" and the Oslo-dummy. In the fifth auction, we include bidder fixed effects. Our preferred model is Model 5.

We observe that the estimated coefficient for the variable "Auction number" is statistically significant in all ten regressions. Our preferred model, Model 5, which includes bidder fixed effects yields an estimated coefficient of -0.766. The economic interpretation is that for each additional auction in which the bidder participates the size of the unit decreases by 0.77 of a square meter. (One square meter is 10.76 square feet.) It is economically meaningful that bidders seek smaller units each time they lose an auction. The reason why is that a smaller unit is less expensive, everything else being the same, and bidding for a less expensive unit increases a bidder's chance of winning with the same bid. Because all ten regressions show statistically significant estimates of "Auction number" and the estimate for Winners (Panel A) of Model 5 is economically meaningful, we reject the null hypothesis of no change in the type of unit for which repeat bidders bid. The effect is smaller for bidders who are observed to lose all auctions. Nevertheless, the estimated coefficient is negative, even if it is statistically insignificant. Thus, we claim that bidders that eventually win an auction tend to seek out smaller units as they participate in more auctions. We reject the null hypothesis of no change.

6.2 A repeat bidder's choice of units with respect to ask price

In Table 5, we repeat the ten regressions for which we reported results in Table 4 using the ask price relative to the annual average in the area for each bidder-auction combination as dependent variable.

We see that the estimated coefficient of "Auction number", i.e., the first row of results, is statistically significant for all five regressions in Panel A. For our preferred model, Model 5, the estimated coefficient of Model 5 is -0.008, which means that for each auction in which a bidder participates the ask price decreases by 0.008 compared to the annual average in the area. In other words, for each auction a bidder loses, the bidder seeks a unit next time that is about one percent cheaper, which is economically meaningful. After all, house prices are often multiples of household income, so one percent represents a substantial monetary

Table 5: Ask price. Regression of the unit ask price on auction number. Partitioned into bidders that lost the last auction and bidders that won the last auction

Panel A: Winners

	(1)	(2)	(3)	(4)	(5)
Auction number	-0.014^{***}	-0.025^{***}	-0.024^{***}	-0.028^{***}	-0.008^{***}
	(0.004)	(0.005)	(0.005)	(0.006)	(0.002)
Auction number*Oslo				0.011	
				(0.008)	
N auctions		✓	✓	✓	
Weeks from first observed		✓	✓	✓	✓
Oslo			✓	✓	✓
Bidder FE					✓
Observations	15,672	15,672	15,672	15,672	15,672
Adjusted R^2	0.001	0.008	0.011	0.011	0.864

Panel B: Losers

	(1)	(2)	(3)	(4)	(5)
Auction number	-0.003	-0.014^{***}	-0.011^{***}	-0.008	0.004^{**}
	(0.003)	(0.004)	(0.004)	(0.006)	(0.002)
Auction number*Oslo				-0.007	
				(0.007)	
N auctions		✓	✓	✓	
Weeks from first observed		✓	✓	✓	✓
Oslo			✓	✓	✓
Bidder FE					✓
Observations	19,532	19,532	19,532	19,532	19,532
Adjusted R^2	-0.00002	0.007	0.017	0.017	0.850

Notes: The table shows results of regressing the ask price relative to the annual average ask price on the auction number for the bidder-auction combination *ba*. The annual average ask price is calculated conditional on whether the housing units are located within Oslo. Panel A gives the results from the subset of bidders that win the last auction they participate in, and panel B gives the result for those not winning the last auction. If bidders are observed to win but still place bids afterwards, they are dropped from the sample. The dummy *N auctions* is a control variable for the number of auctions in which each bidder has participated, *Weeks from first observed* represents how many weeks have passed since the bidder is first observed, *Oslo* indicates whether the auction is located in Oslo, and *Bidder FE* are bidder fixed-effects. We retain bidders that have participated in five or less auctions. Heteroskedasticity-robust standard errors are given in parenthesis. Significance: * p<0.1, ** p<0.05, *** p<0.01.

amount.

The coefficient for “Auction number” is not statistically significant for bidders that never win; for which we use the short notation “losers”. They are represented in Panel B. We observe that for Panel B, the variable “Auction number” is statistically insignificant in our preferred model, Model 5. Moreover, the sign is positive. The interpretation is that this segment of bidders do not change behavior with respect to the ask price of the unit and at least not in the direction that would help win an auction. Potentially, this finding could be part of an explanation for why these bidders do not win any auctions.

The lesson is that successful bidders seek out units that have a lower ask price while unsuccessful bidders do not. We reject the null hypothesis of no change.

6.3 A repeat bidder's highest bid across auctions

Table 6 tabulates results from regressions of the highest within-auction bids. In panel A, we observe mixed results for the regression estimates of the variable of most interest, "Auction number". Two of the five estimates are statistically significant, but the estimate in our preferred model, Model 5, is not statistically significant. Three estimates are negative; two are positive. The estimated coefficient in Model 5 is 1,468, which is a small value when we compare it to the unit values. From Table 3, we recall the the overall mean ask price is 2,85 million. Thus, 1,468 is a modest number compared to the overall mean ask price.

In panel B, two estimates are statistically significant. The estimate of "Auction number" in Model 5 is 27,738, which is substantially larger than the 1,468 found in Panel A. The interpretation is that there is a noticeable difference in detected behavior between winners and losers, since winners stick to their original budget while the losers increase their bids.

This is counter-intuitive at first blush. If losers increase their bids across auctions and winners stick to their budget, intuition would say that it would increase the probability of winning an auction for the first group. However, that only holds as long as the two groups seek out similarly priced units. In the two regressions above, we have found that winners tend to participate in auctions in which the units are smaller and less expensive as they participate in more auctions. What is the winning strategy between increasing bids without changing unit types or changing unit types without changing bids, depends on the strengths of the effects. In the next two regressions, we shall inspect the total effect.

Table 6: Bid. Regression of bidders' highest bids within auctions on auction number. Partitioned into bidders that lost the last auction and bidders that won the last auction

Panel A: Winners

	(1)	(2)	(3)	(4)	(5)
Auction number	40,445.980***	−1,180.146	−18,727.960	−35,412.460*	1,468.140
	(14,048.450)	(17,286.420)	(16,675.670)	(19,082.050)	(6,397.293)
Auction number*Oslo				55,149.420*	
				(29,011.050)	
N auctions		✓	✓	✓	
Weeks from first observed		✓	✓	✓	✓
Oslo			✓	✓	✓
Bidder FE					✓
Observations	15,674	15,674	15,674	15,674	15,674
Adjusted R^2	0.0004	0.002	0.063	0.063	0.900

Panel B: Losers

	(1)	(2)	(3)	(4)	(5)
Auction number	72,917.850***	17,264.980	−2,170.666	5,143.115	27,738.060***
	(11,958.530)	(14,934.260)	(14,715.140)	(18,352.870)	(5,784.857)
Auction number*Oslo				−17,297.150	
				(24,028.830)	
N auctions		✓	✓	✓	
Weeks from first observed		✓	✓	✓	✓
Oslo			✓	✓	✓
Bidder FE					✓
Observations	19,533	19,533	19,533	19,533	19,533
Adjusted R^2	0.002	0.004	0.044	0.044	0.881

Notes: The table shows results of regressing bidders' highest bids within auctions on the auction number. Panel A gives the results from the subset of bidders that win the last auction they participate in, and panel B gives the result for those not winning the last auction. If bidders are observed to win but still place bids afterwards, they are dropped from the sample. The dummy *N auctions* is a control variable for the number of auctions in which each bidder has participated, *Weeks from first observed* represents how many weeks have passed since the bidder is first observed, *Oslo* indicates whether the auction is located in Oslo, and *Bidder FE* are bidder fixed-effects. We retain bidders that have participated in five or less auctions. Heteroskedasticity-robust standard errors are given in parenthesis. Significance: * p<0.1, ** p<0.05, *** p<0.01.

6.4 A repeat bidder's bid levels compared to ask price

Table 7 tabulates results from ten regressions in which we regress the bid-ask price ratio on "Auction number" for bidder-auction combinations *ba*. We observe that all ten regressions yield a statistically significant estimated coefficient for "Auction number". Our preferred model is Model 5, which includes bidder fixed effects. The estimated auction number coefficient for winners is 0.011 and it is statistically significant with a p-value less than one percent. The result in panel B is about half that, 0.006, and statistically significant at the same level.

The interpretation is that for each auction in which an eventually-winning bidder participates the bid-ask price ratio increases by 0.011, or 1.1 percent. This is an economically significant number. Thus, we reject the null of no change in bidding relative to ask price.

This effect also amounts to an answer to the question we posed above of the total effects from either keeping bids but changing types or changing bids but keeping types. Winners are seen to seek out units with ask prices that are sufficiently low that the bids-ask ratio increases. In the next regression, we inspect what is the effect on the competitiveness.

6.5 A repeat bidder's bid levels compared to competing bids

In Table 8, we present ten regressions of the highest bid bidder b makes in auction number a less the minimum bid in auction a divided by the difference between the maximum and minimum bids in auction a. These numbers represent the competitiveness of the bids. We observe that eight of the ten regressions yield statistically significant coefficient estimates for "Auction number". Our preferred model, Model 5, which includes bidder fixed effects, yields estimated coefficients of "Auction number" that are statistically significant for both winners and losers. The estimated coefficients are, respectively, 0.036 and 0.013. The interpretation is that we detect a change in how bidders extend bids relative to competitors as their experience grows. Bidders that win their last auction change more in each auction than do bidders that do not win any auctions. The estimated coefficient of winners is almost three times the estimated coefficient for losers. We reject the null.

In summary, winning bidders tend to stick to their original budget but seek out units that are smaller and less expensive. This strategy implies that as they participate in more auctions, and continue to seek smaller and cheaper units, the competitiveness of their bids increase. In fact, in the range spanned by the minimum bid in an auction to the maximum bid in an auction, winners climb about four points per auction in which they participate.

Table 7: Bids relative to ask price. Regressions of the bid-ask price ratio on auction number. Partitioned into bidders that lost the last auction and bidders that won the last auction

Panel A: Winners

	(1)	(2)	(3)	(4)	(5)
Auction number	0.011***	0.015***	0.015***	0.014***	0.011***
	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)
Auction number*Oslo				0.0005	
				(0.002)	
N auctions		✓	✓	✓	
Weeks from first observed		✓	✓	✓	✓
Oslo			✓	✓	✓
Bidder FE					✓
Observations	15,674	15,674	15,674	15,674	15,674
Adjusted R^2	0.005	0.009	0.017	0.017	0.260

Panel B: Losers

	(1)	(2)	(3)	(4)	(5)
Auction number	0.004***	0.009***	0.009***	0.009***	0.006***
	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)
Auction number*Oslo				−0.001	
				(0.002)	
N auctions		✓	✓	✓	
Weeks from first observed		✓	✓	✓	✓
Oslo			✓	✓	✓
Bidder FE					✓
Observations	19,533	19,533	19,533	19,533	19,533
Adjusted R^2	0.001	0.003	0.009	0.009	0.280

Notes: The table shows results of regressing the max bid relative to the ask price on the auction number for the bidder-auction combination *ba*. Panel A gives the results from the subset of bidders that win the last auction in which they participate, and panel B gives the result for those not winning the last auction. If bidders are observed to win but still place bids afterwards, they are dropped from the sample. The dummy *N auctions* is a control variable for the number of auctions in which each bidder has participated, *Weeks from first observed* represents how many weeks have passed since the bidder is first observed, *Oslo* indicates whether the auction is located in Oslo, and *Bidder FE* are bidder fixed-effects. We retain bidders that have participated in five or less auctions. Heteroskedasticity-robust standard errors are given in parenthesis. Significance: * p<0.1, ** p<0.05, *** p<0.01.

Table 8: Relative bid position of repeat bidders. Regressions of the bid's position along the max-min range on auction number. Partitioned into bidders that lost the last auction and bidders that won the last auction

Panel A: Winners

	(1)	(2)	(3)	(4)	(5)
Auction number	0.038***	0.033**	0.033**	0.016	0.036**
	(0.012)	(0.015)	(0.015)	(0.018)	(0.016)
Auction number*Oslo				0.049**	
				(0.024)	
N auctions		✓	✓	✓	
Weeks from first observed		✓	✓	✓	✓
Oslo			✓	✓	✓
Bidder FE					✓
Observations	1,777	1,777	1,777	1,777	1,777
Adjusted R^2	0.005	0.004	0.003	0.005	0.069

Panel B: Losers

	(1)	(2)	(3)	(4)	(5)
Auction number	0.009**	0.018***	0.019***	0.014**	0.013***
	(0.004)	(0.005)	(0.005)	(0.006)	(0.005)
Auction number*Oslo				0.012	
				(0.008)	
N auctions		✓	✓	✓	
Weeks from first observed		✓	✓	✓	✓
Oslo			✓	✓	✓
Bidder FE					✓
Observations	19,515	19,515	19,515	19,515	19,515
Adjusted R^2	0.0003	0.003	0.003	0.003	0.100

Notes: The table shows results of regressing the relative bid position on the auction number for the bidder-auction combination *ba*. Panel A gives the results from the subset of bidders that win the last auction in which they participate, and panel B gives the result for those not winning the last auction. If bidders are observed to win but still place bids afterwards, they are dropped from the sample. The relative bid position is defined as the ratio between a given bidder b's highest bid in auction a less the global minimum bid in auction a on the difference between the global maximum bid in auction a and the global minimum bid in auction a, $(\text{Highest bid}_{b,a} - \text{Minimum bid}_a)/(\text{Maximum bid}_a - \text{Minimum bid}_a)$. The dummy *N auctions* is a control variable for the number of auctions in which each bidder has participated, *Weeks from first observed* represents how many weeks have passed since the bidder is first observed, *Oslo* indicates whether the auction is located in Oslo, and *Bidder FE* are bidder fixed-effects. The different number of observations in the two panels are due to the construction of the winners. A winner is defined as winning the last unit the bidder is observed in. Therefore, we exclude the auctions the winners are winning, making it also necessary to exclude the bidders only observed twice (the majority in our repeat bidders sample). We retain bidders that have participated in five or less auctions. Heteroskedasticity-robust standard errors are given in parenthesis. Significance: * p<0.1, ** p<0.05, *** p<0.01.

7 Discussion

7.1 Sensitivity

We examine how sensitive our results are to a different specification that includes "Owner-occupied" as a control variable. In Norway, units are classified as self-owned or coops. Generally, households do not differentiate between the two groups, since they are functionally the same. A coop is a unit in which you buy the right to live in a specific unit, and there are few restrictions on how you may re-design the unit. The main difference is that a coop usually has common debt, oftentimes substantial, as the members of the coop collectively finance upgrades. Since common debt is frequently observed for coops, we seek to investigate how it affects the results. Common debt puts restrictions on an individual household's ability to negotiate with the bank and change terms with respect to variable interest versus fixed interest and the length of amortization. Thus, it is fathomable that the difference between a self-owned unit and a coop could affect bidding behavior.

In Table 9, we report results for our preferred Model 5 for all five dependent variables, size, ask price, bid level, bid relative to ask, and relative bid position. We observe that the coefficient estimate for "Auction number" is statistically significant for four dependent variables for winners, but statistically insignificant for one variable, bid level (iii). We are particularly interested in the relative bid position, and we observe that the resulting coefficient estimate is identical to the estimate with the dummy for "Owner-occupied" reported above in Table 8, i.e. 0.036.

For losers, the estimated coefficient is not statistically significant for size and ask price, but statistically significant for bid level, bid over ask, and relative bid position. Again, the estimated coefficient for relative bid position (v) is identical to the estimate without the dummy for "Owner-occupied" reported above in Table 8, i.e. 0.013.

Overall, the patterns are not sensitive to including the control variable "Owner-occupied".

In Table 10, we examine the sensitivity to using year-by-quarter fixed effects instead of

Table 9: Sensitivity to controlling for ownership

Panel A: Winners

	Size	Ask	Bid	Bid over ask	Rel. bid pos.
	(i)	(ii)	(iii)	(iv)	(v)
Auction number	-0.769^{**}	-0.008^{***}	2,516.264	0.011^{***}	0.036^{**}
	(0.342)	(0.002)	(6,404.735)	(0.001)	(0.016)
Owner-occupier	✓	✓	✓	✓	✓
Weeks from first observed	✓	✓	✓	✓	✓
Oslo	✓	✓	✓	✓	✓
Bidder FE	✓	✓	✓	✓	✓
Observations	12,663	15,672	15,674	15,674	1,777
Adjusted R^2	0.843	0.865	0.900	0.261	0.068

Panel B: Losers

	Size	Ask	Bid	Bid over ask	Rel. bid pos.
	(i)	(ii)	(iii)	(iv)	(v)
Auction number	-0.076	0.004^{**}	$28{,}006.170^{***}$	0.006^{***}	0.013^{**}
	(0.262)	(0.002)	(5,793.814)	(0.001)	(0.005)
Owner-occupier	✓	✓	✓	✓	✓
Weeks from first observed	✓	✓	✓	✓	✓
Oslo	✓	✓	✓	✓	✓
Bidder FE	✓	✓	✓	✓	✓
Observations	16,190	19,532	19,533	19,533	19,515
Adjusted R^2	0.849	0.851	0.881	0.281	0.100

Notes: The table shows results of Model 5 (bidder fixed-effects) when controlling for ownership type, by including an indicator being one for *owner-occupiers* else zero. Panel A gives the results from the subset of bidders that win the last auction they participate in, and panel B gives the result for those not winning the last auction. When dependent variable is the unit size, we choose to drop bidders if they are observed to bid on different unit types, e.g. first a detached house and then an apartment. If bidders are observed to win but still place bids afterwards, they are dropped from the sample. The dummy controls *Weeks from first observed* represents how many weeks have passed since the bidder is first observed, *Oslo* indicates whether the auction is located in Oslo, and *Bidder FE* are bidder fixed-effects. For the relative bid position, we drop the winning auction for the winners and winners observed twice, as described in the notes for Table 8. We retain bidders that have participated in five or less auctions. Heteroskedasticity-robust standard errors are given in parenthesis. Significance: * p<0.1, ** p<0.05, *** p<0.01.

"Weeks from first observed". By inspecting the magnitudes and significance of the estimated coefficients for "Auction number" in Table 10, we see that the pattern is intact. In particular, we observe that the estimated coefficients for relative bid position for the variable "Auction number" are 0.038 and 0.017, for winners and losers, respectively. These estimates are very close to the ones reported above using "Weeks from first observed", which were 0.036 and 0.013.

Table 10: Sensitivity to controlling for year-by-quarter fixed-effects

Panel A: Winners

	Size (i)	Ask (ii)	Bid (iii)	Bid over ask (iv)	Rel. bid pos. (v)
Auction number	-0.785^{**}	-0.007^{***}	543.823	0.012^{***}	0.038^{**}
	(0.324)	(0.002)	(6,497.124)	(0.001)	(0.016)
Weeks from first observed					
Oslo	✓	✓	✓	✓	✓
Bidder FE	✓	✓	✓	✓	✓
Year-by-quarter FE	✓	✓	✓	✓	✓
Observations	12,663	15,672	15,674	15,674	1,777
Adjusted R^2	0.843	0.865	0.900	0.265	0.071

Panel B: Losers

	Size (i)	Ask (ii)	Bid (iii)	Bid over ask (iv)	Rel. bid pos. (v)
Auction number	-0.045	0.005^{***}	$26{,}505.690^{***}$	0.007^{***}	0.017^{***}
	(0.250)	(0.002)	(5,836.395)	(0.001)	(0.005)
Weeks from first observed					
Oslo	✓	✓	✓	✓	✓
Bidder FE	✓	✓	✓	✓	✓
Year-by-quarter FE	✓	✓	✓	✓	✓
Observations	16,190	19,532	19,533	19,533	19,515
Adjusted R^2	0.849	0.851	0.881	0.281	0.099

Notes: The table shows results of Model 5 (bidder fixed-effects) when controlling for year-by-quarter fixed effects instead of *weeks from first observed.* Panel A gives the results from the subset of bidders that win the last auction they participate in, and panel B gives the result for those not winning the last auction. When dependent variable is the unit size, we choose to drop bidders if they are observed to bid on different unit types, e.g. first a detached house and then an apartment. If bidders are observed to win but still place bids afterwards, they are dropped from the sample. The dummy controls *Weeks from first observed* represents how many weeks have passed since the bidder is first observed, *Oslo* indicates whether the auction is located in Oslo, *Bidder FE* are bidder fixed-effects and *Year-by-quarter FE* are year-by-quarter fixed-effects. For the relative bid position, we drop the winning auction for the winners and winners observed twice, as described in the notes for Table 8. We retain bidders that have participated in five or less auctions. Heteroskedasticity-robust standard errors are given in parenthesis. Significance: * p<0.1, ** p<0.05, *** p<0.01.

7.2 Robustness

In Table 11, we examine the robustness of using appraisal values instead of ask prices. The main reason is that it is possible to fathom a strategic element in ask prices (Anundsen et al. (2020)). We do have a sub-set of observations with appraisal values, although this sub-sample is limited in time and space. We have replicated the estimates for both the ask price regressions and the bids relative to ask price regressions above. By inspecting the magnitudes and significance of the estimated coefficients for "Auction number", for both winners and losers, in Table 11, we see that the main pattern is intact. In particular, we observe that for winners, the estimated coefficients for 'Auction number" in regressions (i) and (iii) are -0.010 and 0.019 compared to the estimated coefficients in regressions (ii) and (iv) of -0.010 and 0.015. In short, estimates using appraisal values are very close to the ones reported above using ask prices.

Table 11: Robustness for substituting ask prices with appraisal values

Panel A: Winners

	Appraisal (i)	Ask (ii)	Bid over appraisal (iii)	Bid over ask (iv)
Auction number	-0.010^{**}	-0.010^{**}	0.019^{***}	0.015^{***}
	(0.005)	(0.005)	(0.005)	(0.003)
Weeks from first observed	✓	✓	✓	✓
Oslo	✓	✓	✓	✓
Bidder FE	✓	✓	✓	✓
Observations	4,624	4,624	4,624	4,624
Adjusted R^2	0.870	0.869	0.138	0.291

Panel B: Losers

	Appraisal (i)	Ask (ii)	Bid over appraisal (iii)	Bid over ask (iv)
Auction number	0.004	0.004	0.006	0.005^{*}
	(0.004)	(0.004)	(0.004)	(0.003)
Weeks from first observed	✓	✓	✓	✓
Oslo	✓	✓	✓	✓
Bidder FE	✓	✓	✓	✓
Observations	6,495	6,495	6,495	6,495
Adjusted R^2	0.846	0.846	0.130	0.245

Notes: The table shows results of model (5) (bidder fixed-effects) when substituting ask prices with appraisal values. Column (1) shows the result from when the dependent variable is the appraisal value relative to the annual average appraisal value with Oslo and non-Oslo. Column (2) supplements the first column with a replication of Model 5 in Table 5 but for the exact same sample as in column (1). Columns (3) and (4) repeats the same exercise for Model 5 in Table 7. *Weeks from first observed* represents how many weeks have passed since the bidder is first observed, *Oslo* indicates whether the auction is located in Oslo, and *Bidder FE* are bidder fixed-effects. We retain bidders that have participated in five or less auctions Heteroskedasticity-robust standard errors are given in parenthesis. Significance: * p<0.1, ** p<0.05, *** p<0.01.

8 Concluding remarks and implications

We present patterns that emerge when we study a novel data set on Norwegian bidding logs in housing auctions. The data set is sourced from Norway's second largest realtor firm, DNB Eiendom, and contains more than one million bids and about 200,000 housing transactions.

We find that repeat bidders seem to change behavior in respect to size and ask price of the units for which they bid. As they participate in more auctions, i.e., as they learn more, they tend to bid for units that are smaller and less expensive. These effects are much stronger for bidders that win their last auction compared to bidders that never win an auction. However, bidders that win their last auction do not seem to change the level of their bids. Instead, they tend to stay within budget across auctions. In contrast, the bidders that never win has a small inclination to increase their bids as they participate in more auctions. As a result of the combined effects of reductions in size and ask price while maintaining budget discipline, winning bidders are observed to increase their bids relative to ask price and increase their bids relative to competing bids. For each auction, winning bidders climb 3.6 points on the auction range from the auction's minimum bid to the auction's maximum bid. In contrast, losing bidders climb only 1.3 points.

The implications are that bidders win when they are able to position themselves favorably compared to competitors while staying within budget. Thus, bidders that keep losing auctions could change the trend by accepting that their budget does not allow the units for which they compete. Instead, they would be wise to seek out somewhat smaller and less expensive units.

References

Anundsen, A. K., Larsen, E. R., and Sommervoll, D. E. (2020). Strategic price-setting and incentives in the housing market. *Housing Lab working papers*, 1.

Anundsen, A. K., Nenov, P., Larsen, E. R., and Sommervoll, D. E. (2022). Pricing and incentives in the housing market. *Housing Lab working papers*, 3.

Arefeva, A. (2017). How Auctions Amplify House-Price Fluctuations. Johns Hopkins Carey Business School Research Paper 17-12, Available at SSRN 2688859.

Arefeva, A. and Meng, D. (2020). How to set a deadline for auctioning a house. *Available at SSRN 3520197*.

Avery, C. (1998). Strategic jump bidding in english auctions. *The Review of Economic Studies*, 65(2):185–210.

Bikhchandani, S. and Riley, J. G. (1991). Equilibria in open common value auctions. *Journal of Economic Theory*, 53(1):101–130.

Börgers, T. and Dustman, C. (2005). Strange bids: Bidding behaviour in the united kingdom's third generation spectrum auction. *The Economic Journal*, 115:551–578.

Bucchianeri, G. W. and Minson, J. A. (2013). A homeowner's dilemma: Anchoring in residential real estate transactions. *The Review of Economic Behaviour and Organization*, 89:235–260.

Chow, Y. L., Hafalir, I. E., and Yavas, A. (2015). Auction versus negotiated sale: evidence from real estate sales. *Real Estate Economics*, 43(2):432–470.

Dodonova, A. (2017). Preemptive bidding and pareto efficiency in takeover auctions. *Economics Letters*, 159:214–216.

Ettinger, D. and Michelucci, F. (2016a). Creating a winner's curse via jump bids. *Review of Economic Design*, 20:173–186.

Ettinger, D. and Michelucci, F. (2016b). Hiding information in open auctions with jump bids. *The Economic Journal*, 126(594):1484–1502.

Ettinger, D. and Michelucci, F. (2019). Manipulating information revelation with reserve prices. *Annals of Economics and Statistics*, 133:87–92.

Genesove, D. and Hansen, J. (2023). Auctions and negotiations in housing price dynamics. *Review of Economics and Statistics*, pages 1–39.

Gilbukh, S. and Goldsmith-Pinkham, P. (2023). Heterogeneous real estate agents and the housing cycle. Technical report, National Bureau of Economic Research.

Han, L. and Strange, W. C. (2014). Bidding wars for houses. *Real Estate Economics*, 42(1):1–32.

Horowitz, J. L. (1986). Bidding models of housing markets. *Journal of Urban Economics*, 20(2):168–190.

Hungria-Gunnelin, R. (2018). An analysis of auction strategies in apartment sales. *Journal of European Real Estate Research*.

Isaac, R. M., Salmon, T. C., and Zillante, A. (2007). A theory of jump bidding in ascending auctions. *Journal of Economic Behavior & Organization*, 62(1):144–164.

Levin, E. J. and Pryce, G. B. (2007). A statistical explanation for extreme bids in the house market. *Urban Studies*, 44(12):2339–2355.

Mateen, H., Qian, F., and Zhang, Y. (2021). The microstructure of the us housing market: evidence from millions of bargaining interactions. *Available at SSRN 3727150*.

Merlo, A. and Ortalo-Magne, F. (2004). Bargaining over residential real estate: evidence from england. *Journal of urban economics*, 56(2):192–216.

Milgrom, P. R. and Weber, R. J. (1982). A theory of auctions and competitive bidding. *Econometrica: Journal of the Econometric Society*, pages 1089–1122.

Smith, E. (2020). High and low activity spells in housing markets. *Review of Economic Dynamics*, 36:1–28.

Sommervoll, D. E. (2020). Jump bids in real estate auctions. *Journal of Housing Economics*, 49:101713.

Sønstebø, O. J., Olaussen, J. O., and Oust, A. (2021). Opening bid strategies in english auctions. *Journal of Real Estate Research*, 43(1):123–143.