

February 2021

“Imperfect Information, Learning and
Housing Market Dynamics”

Christophe Alain Bruneel-Zupanc

Imperfect Information, Learning and Housing Market Dynamics*

Christophe BRUNEEL-ZUPANC[†]

Toulouse School of Economics

January 24, 2020

Abstract

This paper examines the decision problem of a homeowner who maximizes her expected profit from the sale of her property when market conditions are uncertain. Using a large dataset of real estate transactions in Pennsylvania between 2011 and 2014, I verify several stylized facts about the housing market. I develop a dynamic search model of the home-selling problem in which the homeowner learns about demand in a Bayesian way. I estimate the model and find that learning, especially the downward adjustment of the beliefs of sellers facing low demand, explains some of the key features of the housing data, such as the decreasing list price overtime and time on the market. By comparing with a perfect information benchmark, I derive an unexpected result: the value of information is not always positive. Indeed, an imperfectly informed seller facing low demand can obtain a better outcome than her perfectly informed counterpart thanks to a delusively stronger bargaining position.

JEL classification: D83, R2, R3.

Keywords: housing, pricing, imperfect information, Bayesian learning.

*This work was supported by the Agence Nationale de la Recherche - ECHOPPE project, grant ANR-17-CE41-0008-01. I gratefully acknowledge my advisor Thierry Magnac for his regular guidance and feedback. I also thank Robert Miller, John Rust, Michel Le Breton and the participants of the TSE Econometrics and Applied Microeconomics workshops and the CMU Micro Lunch for their comments, plus the participants of the IAAE and ESEM 2018 conferences.

[†]Email: christophe.bruneel@gmail.com

1 Introduction

Real estate transactions involve large financial amounts (\$206k, on average, in Pennsylvania between 2011 and 2014) and take time (109 days). A lengthy or negative outcome of the home-selling process can make a substantial difference to the seller's well-being, and making a successful sale can be challenging. Contrary to many other markets, the homeowner does not post a take-it or leave-it price. Instead, she posts a list price which impacts transaction outcomes less directly (Han and Strange, 2016). Indeed, even though the list price influences negotiations between sellers and buyers, the transaction frequently occurs at a different sale price (85% of the time in my sample).

Of utmost importance to the seller is that a higher list price should yield a higher sale price but at the expense of a longer time on the market (Miller and Sklarz, 1987). This trade-off is complex, especially because the seller has imperfect information about the demand (Salant, 1991). Indeed, houses are highly differentiated assets and the seller is often unable to observe more than a few recent transactions of similar properties nearby. This lack of information impacts the seller's welfare (Anenberg, 2016).

In this article, I investigate the home-selling decision of imperfectly informed sellers. I build a single-agent dynamic search model of the housing market in which a rational seller is uncertain and learns about demand in a Bayesian way. I estimate the model using an original dataset of real estate transactions in Pennsylvania between 2011 and 2014. The estimated model yields insights on the optimal list pricing strategies and how information frictions affect them. In particular, I estimate the cost of uncertainty for the seller.

The literature establishes several stylized facts about the home-selling problem, and my model aims to explain three of them. First, we observe that transactions are occurring at a price mostly below, but also sometimes above or exactly equal to, the list price (Merlo and Ortalo-Magne, 2004). Second, the housing market is illiquid. Third, the list prices are duration-dependent and adjust downwards throughout the listing process, even when market conditions seem stable (Salant, 1991).

My main contribution is to make sense of the two first facts by modeling the number of buyers on the market explicitly, independently from their valuations for the house.

First, there can be several buyers on the market. This allows me to explain the relationship between the list and the sale prices. Indeed, as explained earlier, the sale price is determined after bargaining between the seller and the buyer(s). In these negotiations, the list price serves as a ceiling: the seller commits to accept any offer equal to or above it (Horowitz, 1992; Chen and Rosenthal, 1996a,b). Thus, if there is only one buyer for the house, he will never make an offer greater than the list price. Sales will occur either below, or exactly at, the list price (Arnold, 1999), but never above it. I endogenize

such above list prices by introducing competition between several buyers, as in [Han and Strange \(2016\)](#). More precisely, I derive the solution to a sequential bargaining game à la [Rubinstein \(1982\)](#) with several buyers and one seller who commits to accepting offers higher than her list price. In this setup, the seller can obtain a price above her list price. Indeed, when she bargains with the highest valuation buyer, the second-highest buyer's valuation serves as an outside option for the seller ([Shaked and Sutton, 1984](#)). This second-highest valuation can be above the list price, which sometimes allows the seller to obtain a price above the ceiling set by the list price. Through this mechanism, modeling the number of buyers allows us to endogenize sales above the list price in order to gain a better understanding of the phenomenon. This is especially relevant since the proportion of sales greater than the list price has been non-negligible in recent years: from only 4% in the mid-1990s ([Han and Strange, 2014](#)) to 15% in the US during the 2000s boom, 10% after the bust ([Han and Strange, 2016](#)) and 11.42% in my sample.

Second, modeling the number of buyers in addition to their valuations helps to understand the market's illiquidity and how much of it is caused by the list price. Indeed, the list price serves as a signal to attract buyers. The higher this price, the more buyers expect to pay and the less likely they are to visit the house and bargain to buy it. If we observe no transaction, it is either because there are no buyers on the market, or because the list price is too high and repels buyers. My specification of the demand helps to disentangle the 'list price induced illiquidity' (low buyers valuations) from demand illiquidity (no buyer on the market). For example, even with an extremely low list price (such that even low valuation individuals would be willing to buy), the seller is not guaranteed to sell her home, as there may simply be no buyer on the market during this period. My model explains this kind of illiquid demand situation observed in the data. Moreover, using data on time on the market, list and sale prices, I can identify the valuation process independently from the market thinness.

The third key fact is that the list price is non-stationary. To model this, I assume that uncertain sellers dynamically learn about the demand. The gradual acquisition of information about an uncertain demand can explain why a list price varies, even if market conditions are unchanged ([Anenberg, 2016](#)). More precisely, the seller is uncertain about the liquidity process (the number of buyers on the market) and progressively learns about it by observing visits. Recall that a buyer only visits the house if his valuation is high enough with respect to the signal given by the list price. Thus, by only observing visits, the seller does not necessarily observe all of the buyers on the market. If she sets a low list price all buyers will visit and she directly observes the number of buyers on the market. Otherwise, by setting a higher list price she may select some buyers, and learn less about the true number of buyers on the market. In this original application of Bayesian learning, the decision variable (list price) influences the informational flow (by influencing the odds of observable entries) and thus the learning pace. In

addition to the usual trade-off between high price and short time on the market, the list price also embeds a learning externality here: the lower the list price, the faster you learn.

Estimating the parameters of such a model in which optimal strategies are time-dependent (non-stationary) requires detailed microdata. I use an original dataset of approximately 100,000 complete listing histories (dated initial list price and revisions up to the final sale) of sold (with the help of a realtor) single-family homes in Pennsylvania between 2011 and 2014. I collected the data from the American real estate website, *zillow.com*. I study the home-selling problem in this new context (Pennsylvania 2011-2014) and observe similar stylized facts as the one described previously (see [Knight, 2002](#); [Merlo and Ortalo-Magne, 2004](#); [Han and Strange, 2014](#); [Anenberg, 2016](#)).

In terms of results, my learning model (seven structural parameters only) closely matches the data. I find that progressive learning (in particular the downward adjustment of the initial rational beliefs of sellers facing low demand) is key to explaining the observed decline of the list price. The model also reproduces the distribution of time spent on market.

Finally, this paper contributes to the literature on the role of overconfidence ([Odean, 1998](#); [Piazzesi and Schneider, 2009](#)). I estimate the cost of uncertainty or value of information. To do so, I simulate the model and compare it to a perfect information benchmark. Counterintuitively, I find that being misinformed is not necessarily bad for the seller, at least for sellers facing low demand who are thus initially ‘overconfident’. For them, the value of additional information may even be negative, as an overconfident seller may obtain a further discounted sale price than her perfectly informed counterpart. Indeed, by being wrongfully overconfident, she overestimates her reservation value. She has a genuinely stronger position in the bargaining game, resulting in a higher sale price. However, she also refuses offers more easily and sets higher list prices, implying fewer visits, leading to a longer time spent on the market because of the overestimation. There is an ‘overconfidence area’, where the gain in sale price offsets the overly long time spent on the market, resulting in a better outcome. Thus, sellers facing low demand can be better off by not knowing this and starting with imperfect information rather than being perfectly informed.

On the other hand, sellers facing high demand suffer twice from the imperfect information: they do not pick the list price that maximizes their true outcome, and they have a weaker bargaining position.

Related Literature: The role of the asking price has been extensively studied in the literature (see [Han and Strange, 2015](#), for a complete survey).

First, the literature establishes some key stylized facts about the relationship between sale and list prices. On average, the ratio of sale to list price is around 96% ([Case and Shiller, 1988, 2003](#); [Merlo and Ortalo-Magne, 2004](#); [Han and Strange, 2014](#)). Sales frequently occur below the asking price: generally close to 75% of the recent sales in the US ([Carrillo, 2013](#); [Anenberg, 2016](#)). Sometimes transactions occur at a price greater than the list price. As mentioned, this was very uncommon in the 1990s: around 4% of all transactions in England ([Merlo and Ortalo-Magne, 2004](#)) and in the US ([Han and Strange, 2014](#)). However, this has recently become more frequent (15% during the 2000s boom, and is currently around 10%).

There also exists a mass point of sales (from 10% to 25%) occurring exactly at the list price. To explain this, most housing search models treat the list price as a binding ceiling ([Horowitz, 1992](#); [Chen and Rosenthal, 1996a,b](#)).

Yet, as already mentioned, these models cannot explain why some sales occur above the list price. [Han and Strange \(2016\)](#) fill the gap by modeling a bargaining search model between one seller and several buyers. As in my model, competition can generate prices above the list price. In fact, their bargaining rule is a particular distinction of this study: I allow for any valuations (for the seller and buyers) while they specify a discrete distribution with only two types (high and low valuations buyers). They picked this bargaining rule based on common sense to fit the data. I go further and show theoretically that it is the solution to a sequential bargaining problem à la [Rubinstein \(1982\)](#).

The impact of the list price on time on the market has also been studied. A lower list price increases the probability of visits and sale of the house ([Salant, 1991](#); [Horowitz, 1992](#); [Carrillo, 2012](#)). Similarly, within a listing process, a downward adjustment of the list price increases the probability of sale ([de Wit and van der Klaauw, 2013](#)). Using unique survey data about a buyer's search behavior, [Han and Strange \(2016\)](#) show that the list price directs the buyer in his search.

In addition, sellers of atypical houses are more likely to spend a longer time on the market ([Haurin, 1988](#)). This is consistent with the imperfect information story developed in my paper. Indeed, sellers of atypical houses observe fewer past transactions of houses similar to theirs, thus, they are less informed about the demand they face than their neighbors are. This may explain the difference in behavior: as shown in reduced form evidence by [Anenberg \(2016\)](#), information frictions impact the seller's decision.

Fewer studies have focused on the dynamics of list prices. In order to address this, it is necessary to build a model with some time-dependence. Only two recent papers have done this, to date: [Merlo](#)

et al. (2015) and Anenberg (2016).

Merlo et al. (2015) focus on dynamics, with their main objective to explain list price stickiness: 77.3% of the sales occur without any list price adjustment, and 20.8% with only one adjustment. They show that an extremely small menu cost can generate the observed stickiness. Because it is already well explained in their paper, I abstract from price stickiness in my model. To make sense of the optimal decline in list price, they estimate a model with a rich time-dependent arrival probability function. Overall, they fit their data very closely at the expense of a heavily parametrized model, while I get a decent fit using only seven structural parameters by introducing learning and imperfect information in the model.

Anenberg (2016) is the closest paper to mine. He builds a home-selling model with imperfect information and learning. He formulates and estimates a model in which the seller is uncertain about the buyer's valuation for her property. However, she has a prior about the mean of the valuations distribution. In this context, as in my work, the gradual acquisition of information by the seller can explain the time-varying list price choices (which declines, but also rarely increases, for example). Because of information frictions, a short-run aggregate price may take a longer time to adjust.

The main contrast with this paper is that I model the number of buyers explicitly, and not only their valuations for the house. This allows the study of some aspects of the home-selling problems, which were not the focus of Anenberg's paper. First, I can give a micro-foundation to the bargaining side of the problem. I endogenously generate sales above the list prices, while models with only valuations cannot. Anenberg (2016) generates them with an exogenous probability, for instance. Given the important share of sales above the list price (more than 10% of sales in the US after 2010), being able to endogenize and explain them is crucial: even more so because of their pro-cyclicality (Liu et al., 2014). Moreover, my micro-founded bargaining rule also allows me to derive the paradoxical result that uncertainty can be beneficial to the seller. Finally, using data on list prices, sale prices and time on the market, I can separately identify low valuations (list-price induced illiquidity) from illiquidity (no buyer).

This paper proceeds as follows. Section 2 describes the data. Section 3 develops a dynamic micro-search model of the home-selling problem in which sellers learn. Section 4 details the estimation methods and Section 5 presents the estimation results. Section 6 uses the estimated model to analyze the value of information. I conclude the paper in Section 7.

2 Data

2.1 Source

My data contain transaction records of properties sold between August 2011 and July 2014 in Pennsylvania, gathered from the real estate platform, *zillow.com*. The website is one of the leading online real estate marketplace sites. It is a real estate listing aggregator which gathers listings from real estate agents, in partnership with MLS services and from private national companies (Century 21, Coldwell Banker and Sotheby's, for example). Consequently, this is a collection of some of the most exhaustive data about real estate transactions in Pennsylvania.

I focus on the sales of single-family homes for which there is a complete history record of the transaction available. For each of these sales, the data include the usual property attributes (square footage, lot size, number of beds, number of baths, year when the house was built, etc.). For approximately 33% of all transactions, a complete record of the last transaction history is available. As displayed in Figure 1, this history contains all of the seller's decisions: the initial list price, its eventual adjustments through the sale process, potential intermediary listing removals, and the final sale price.

Figure 1: Example of transaction history record

DATE	EVENT	PRICE	\$/SQFT	SOURCE
11/20/13	Sold	\$328,000 -0.6%	\$140	Public Record
11/18/13	Listing removed	\$329,900	\$141	BHHS Fox & Roa...
10/02/13	Price change	\$329,900 -5.7%	\$141	Prudential Fox...
08/29/13	Price change	\$349,900 -5.2%	\$150	Prudential Fox...
08/06/13	Price change	\$369,000 -2.4%	\$158	Prudential Fox...
07/12/13	Price change	\$378,000 -5.5%	\$162	Prudential Fox...
06/08/13	Listed for sale	\$399,900	\$171	Prudential Fox...

This property was sold after five months on market in 2013, after four list price adjustments. It was finally sold at a price of \$328,000, slightly below the final list price of \$329,900.

In order to study list price dynamics and time on the market, I focus only on transactions for which the history is available. It yields a standard selection bias. Indeed, the detailed history is available for properties of better quality, resulting in a significantly higher sale price than the complete pool

of transactions on average. However, this selection bias is present in most data used in the reference literature, which are obtained from real estate agencies. In theory, zillow.com (hereafter, Zillow) allows homeowners to list their house on their own ('for sale by owner'). Unfortunately, these listings represent less than 1% of the observations. As a consequence, these data cannot be used to understand the seller's choice to resort (or not) to using a real estate agent (as modeled by [Salant \(1991\)](#)), and I focus only on transactions in which a real estate agent has intervened (as in [Merlo et al. \(2015\)](#) or [Anenberg \(2016\)](#)).

Zillow's data exhibit two main flaws. First, the data are right-censored: I only observe transactions which ended up in a success (sale). I do not observe properties that are still on the market or were withdrawn by the sellers (except if they were relisted and sold afterwards). In the transaction history of sold properties, I know whether or not the seller previously decided to pull her property off the market (before relisting it for the final sale that I observe). As I observe the complete sale history, I know the price choices/adjustments made before a potential withdrawal. A total of 17% of the final sales have been withdrawn from the market and relisted later before being sold.¹ One could imagine that a seller who has already tried to sell her property (and failed) differs from a 'new' seller. In particular, these two types of sellers should differ in terms of the information they have about the market conditions (one of the two having accumulated information about the market demand with her failed listing). In order to avoid these differences, and since I do not focus on the withdrawal decision here, I drop properties which were withdrawn at least once from my sample and focus only on 'first time' sellers.

Second, as in most of the real estate data, I lack information about buyers' offers and the eventual rejection of these offers by the sellers. Thus, I am forced to keep the buyer's side of my model relatively simple (contrary to [Merlo et al. \(2015\)](#) who have information about buyers' bids and can use it to model their bargaining process, for example).

After selecting the observations and cleaning the data, I have 97,451 real estate transactions in Pennsylvania between August 2011 and July 2014. The data are described in the next section.

¹This implies that the proportion of sellers who fail to sell their property is larger than 17% overall, since I only observe a part of the withdrawals: the properties that were relisted and sold afterwards.

2.2 Summary statistics

Table 1: Summary Statistics

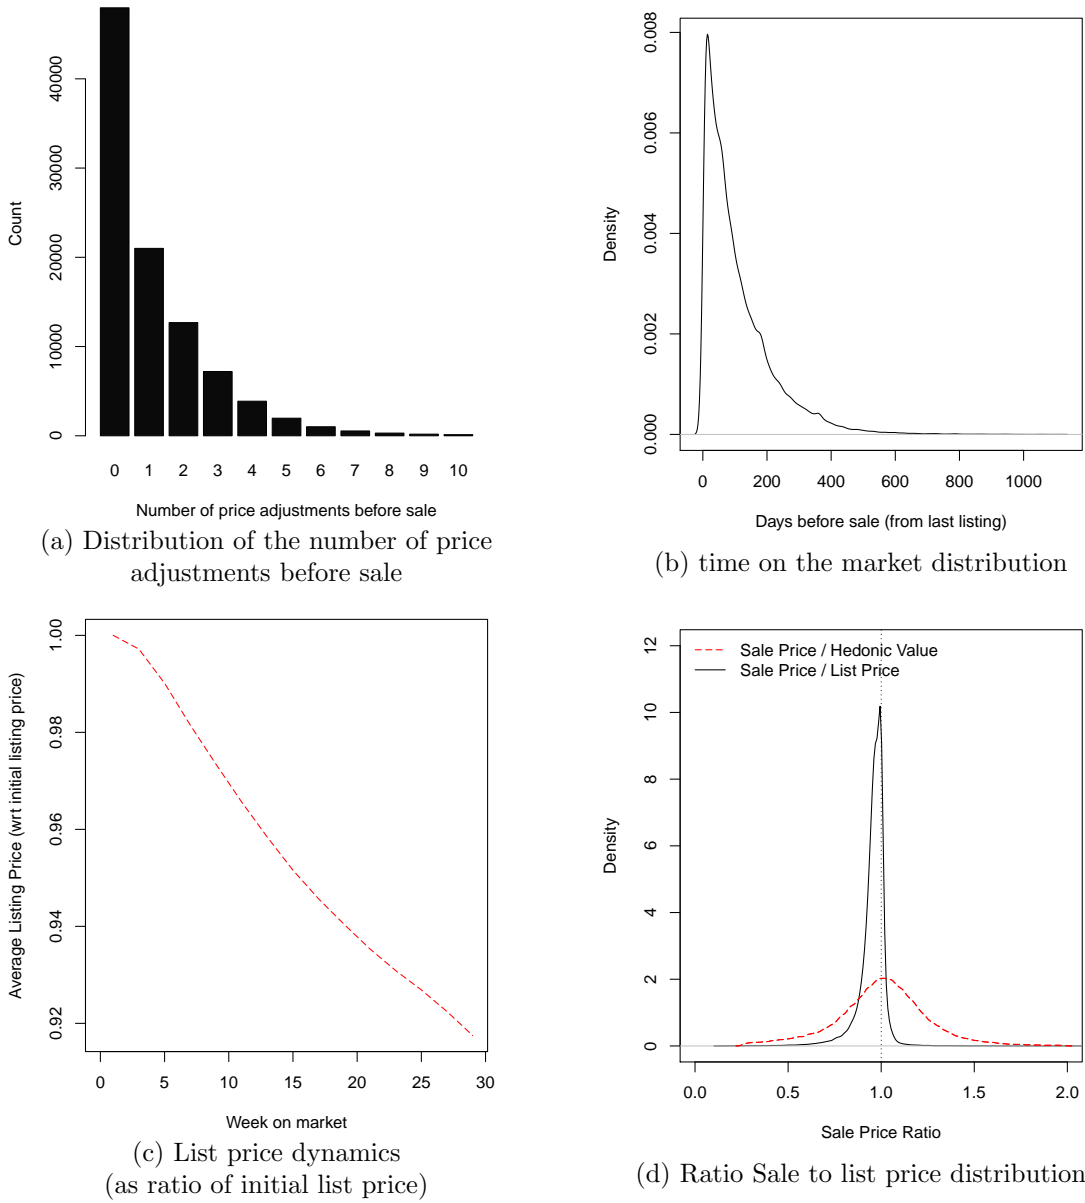
Variable	Mean	Std. Dev.	Min	Median	Max
<i>Price and Timing:</i>					
Price	206512	122209	15000	179000	745000
Days before sale (from last listing)	109	108	0	76	1112
Days before sale (from first listing)	118	110	7	84	1141
Number of adjustment before sale	1.12	1.56	0	1	10
Proportion of listing with adjustments	0.5046				
Ratio sale/Final list price	0.9546	0.0847	0.1316	0.9679	4.3337
Ratio sale/Initial list price	0.9121	0.1089	0.1316	0.9365	4.3337
Proportion of sale price > Final list price	0.1142				
Proportion of sale price = Final list price	0.1495				
<i>Properties characteristics:</i>					
Living area (sqft)	1850	711	780	1682	4572
Number of beds	3.34	0.74	1	3	7
Number of baths	2.1	0.81	0.5	2	5.5
<i>Number of transactions</i>					
	97451				
<i>Number of census tracts</i>					
	1806				

Table 1 and Figure 2 present the summary statistics of my sample. The Pennsylvanian data exhibit common stylized facts to those observed in the literature and which have been extensively detailed by Merlo and Ortalo-Magne (2004). First, prices are sticky and often not adjusted (see Figure 2a). However, in my sample they are not as ‘sticky’ as usual: only 50% of the sellers sell their properties without changing their list price at least once. This percentage is generally closer to 75%: equal to 76.79% in Merlo and Ortalo-Magne (2004) in the UK for example, and also closer to 75% in another Zillow data set of properties on the East Coast of the United States. I choose not to focus on the stickiness and not to model it. It has already been well explained by Merlo et al. (2015) and including a menu cost in my model would be too computationally costly (I would need to add the previous period list price as a state variable) and provide no interesting new insight.

The average time on the market is about 16 weeks (Figure 2b). This is within the usual range: higher than that observed by Merlo et al. (2015) in the UK (about 10 weeks), and slightly lower than that observed by Anenberg (2016) in San Francisco and Los Angeles (about 18 weeks on market).

The list price is decreasing through the selling process: minus 8% between the first and the 30th week (Figure 2c). Most sales happen below the list price (73.63%), many occur exactly at the list price (14.95%) and the 11.42% remaining occur above it. This number of sales above the list price

Figure 2: Descriptive Statistics



is considerably greater than in the English data (3.9% in [Merlo et al., 2015](#)): this is why my model emphasizes the endogenization of sales above the list price.

Unobserved private information: Figure 2d shows the distribution of sale prices normalized either by the final list prices or by the predicted prices (hedonic values). The distribution of prices normalized by the hedonic values of the property is less concentrated than the one normalized by list prices. The list price contains extra private information about the property value, which is not contained in the explanatory variables used in the hedonic estimation (number of bathrooms, number of beds, living area, census tracts, etc.). I find that the list price is a better predictor of the sale price than the hedonic fitted price, as is well known in the literature ([Horowitz, 1992](#); [Merlo et al., 2015](#)). To

soften the impact of this private information unobserved by the econometrician, I use mostly ‘relative moments’ normalized by the list price when I estimate the model, for example, sale price distribution relative to the final list price, or list price dynamics relative to initial list price.

3 Model

I model a discrete-time (with periods of two weeks), infinite horizon problem of a homeowner deciding to sell her property with the objective of maximizing the sale price. As in [Anenberg \(2016\)](#), I introduce uncertainty and Bayesian learning into this framework. To explain the previously described stylized facts, one of the main features of the model is that the seller is imperfectly informed about the true demand. More precisely, she has rational expectations about the distribution of the number of buyers in the population, but she does not know the distribution of buyers interested in her house specifically for each period, and she is learning about it progressively.

Each period t , given her information set, the seller picks a list price p_t^L in order to maximize her expected gains from the sale. The chosen optimal list price will balance a classic trade-off between high sale price and short time on the market. To this classic trade-off, the model adds a learning externality to the list price decision: *ceteris paribus*, a lower list price choice allows the seller to learn more quickly about the market conditions.

In addition to her dynamic list price decision, the seller also decides whether or not to accept an offer, knowing that if she refuses (or if she receives no offer), she will incur a holding cost of keeping her house for sale one more period, modeled as a discount factor δ .

In what follows, I first describe the within-period game: the explicit bargaining rules defining the sale price, the demand modelization and the seller’s learning process. Next, I focus on the seller’s dynamic optimization problem specification.

3.1 Bargaining rules

The sale price is determined after a sequential bargaining game ([Rubinstein, 1982](#)) of offers/counteroffers with complete information between one seller (valuation v^s) and n inspecting/visiting buyer(s) (with ordered valuations $v_{(1)}^b < v_{(2)}^b < \dots < v_{(n)}^b$). The list price p^L serves as a commitment device in the model: if a buyer makes an offer greater or equal to it, the seller has to accept it and sell him the property. This bargaining game between n buyers and one seller with complete information has a unique subgame-perfect equilibrium outcome (see [Appendix A](#)) which is my *bargaining rule*:

- if $v^s > v_{(n)}^b$ (or if $n = 0$): no sale.

- Otherwise, gains from trade exist with at least one buyer, thus under complete information the seller will sell her house to the highest valuation buyer at the price p^S , where

$$p^S = \max\left\{v_{(n-1)}^b, \min\left(p^L, v^s + \frac{1}{2}(v_{(n)}^b - v^s)\right)\right\}$$

If $n = 1$, remove the $v_{(n-1)}^b$ part (or consider it = 0).

In words, trade only occurs if there are gains from trade ($v_{(n)}^b \geq v^s$). If this is the case, under complete information the seller will sell her property to the highest valuation buyer (buyer_(n)) from whom she can extract the higher sale price.

Without the presence of other buyers and without a list price, bilateral sequential bargaining between the seller and buyer_(n) would yield the classic [Rubinstein \(1982\)](#) outcome whereby they share the ‘transaction gains’. The seller gets a portion ϕ of $(v_{(n)}^b - v^s)$ while the buyer gets the rest (the portion $1 - \phi$). Thus the sale price would be $p^S = v^s + \phi(v_{(n)}^b - v^s)$. For simplicity, I make assumptions such that they share the transaction gains equally ($\phi = 0.5$).²

Moreover, the seller has to accept any offer higher or equal to the list price. Thus, if the bilateral bargaining price is greater than p^L , this gives an opportunity for the buyer to make a lower offer, equal to p^L , that the seller must accept. As a consequence, bilateral bargaining between the seller and buyer_(n) with a list price would yield the sale price $p^S = \min\left(p^L, v^s + \frac{1}{2}(v_{(n)}^b - v^s)\right)$.

Now, if another buyer (the second-highest valuation buyer, denoted buyer_(n-1)) is present on the market, his valuation $v_{(n-1)}^b$ can serve as an outside option for the seller.³ If $v_{(n-1)}^b$ is high enough (greater than the bilateral bargaining with list price outcome), the seller can threaten buyer_(n) to sell her property to buyer_(n-1) instead. Competition between the two buyers forces buyer_(n) to offer at least $p^S = v_{(n-1)}^b$ in order to ensure that the seller sells him the property. Adding this competitive outside option to the bilateral bargaining of the seller with only buyer_(n) yields the general sale price formula.

The main advantage of this bargaining rule is that it can endogenously generate prices below the list price ($v^s + 1/2(v_{(n)}^b - v^s)$), at the list price (p^L), and above it (in case of competition, with $v_{(n-1)}^b$ which can be higher than p^L sometimes). The challenge here is to reproduce it in proportions comparable to the data (15% equal to p^L , 11.5% above and the rest below).

The bargaining rule is known by the seller (and the buyers). This implies that, though the seller does not know the exact buyers’ valuations before meeting them, knowing the bargaining rule allows her

²To end up with an equal sharing of the pie as in [Rubinstein \(1982\)](#) sequential bargaining, I assume either that the bargaining period length tends to zero or that all the agents are equally impatient with impatience factor tending to one (see [Appendix A](#)).

An alternative to the equal share would be to give all the bargaining power to the seller (i.e. having her infinitely more patient than the buyer), as done implicitly in [Anenberg \(2016\)](#).

³Notice that only the second-highest valuation can impact the sale price: ceteris paribus the other buyers’ valuations are irrelevant.

to compute what the hypothetical bargaining outcomes may be (sale or not, and eventual sale price) for any scenario. Therefore, for a given number of buyers, if she knows the buyers' valuations distribution, the seller is able to build an expectation about the bargaining outcomes (using order statistics of the highest and second-highest valuation for the given number of buyers). If she also has an idea of the distribution of the number of buyers on the market, she can compute her expected profit from sale as a function of the list price. Then she can pick the list price to maximize it: this is what the seller does in this model.

The demand side of the model is detailed in the next section.

3.2 Buyer

The demand side is split in two main components: the number of buyers on the market during the period ($N_t^{market} \sim Poisson(\lambda)$) and the valuations of each of these buyers for the given property ($v^b \sim \mathcal{LN}(\mu, \sigma)$).

For each property s and for each period t of two weeks, the number of buyers on the market potentially interested in the property (N_t^{market}) is drawn from a $Poisson(\lambda_s)$ distribution. The rate of arrival λ is a key demand parameter in the model. Each seller faces a specific λ_s drawn from the true distribution $Gamma(\alpha_0, \beta_0)$. They have rational expectation and know that $\lambda_s \sim Gamma(\alpha_0, \beta_0)$, but they do not know the value of their individual draw λ_s . With the information they obtain through their own listings, the sellers progressively update their initial rational belief about λ_s using Bayesian learning.

The reservation values of every buyer on the market for a given property are defined as follows:

$$\begin{aligned} V^b &= \eta_s \exp(\theta_b) \\ \iff v^b &:= \frac{V^b}{\eta_s} = \exp(\theta_b) \text{ with } \theta_b \sim \mathcal{N}(\mu, \sigma) \\ \iff v^b &\sim \mathcal{LN}(\mu, \sigma) \end{aligned}$$

where θ_b represents the buyer specific taste for a given property and η_s represents the property intrinsic/predicted value (estimated via hedonic regression). I assume that the buyer knows his taste θ_b for the property, as well as the property intrinsic value η_s . Thus, he also knows his reservation value (V^b) and his reservation value normalized by the hedonic value (v^b). I focus on normalized values (v^b) rather than real monetary values (V^b) in order to compare any type of homes on the same scale. By doing this, I implicitly assume a linear homogeneity of the home-selling problem between the different properties, as in [Merlo et al. \(2015\)](#). In other words, I assume that a \$10,000 deviation for a property worth \$100,000 is perceived similarly to a \$20,000 deviation for a property worth \$200,000. In this way, I

build a single representative problem for every seller, independently of the ‘quality’ of their properties. Then, under the linear homogeneity assumption, I can compare data counterparts to the model price outcomes; that is, the data prices normalized by their hedonic values.

The seller knows the buyer’s reservation value distribution $v^b \sim \mathcal{LN}(\mu, \sigma)$ (but she does not know the buyer’s exact taste shock realization before entering into bargaining with him). Thus, to build expectations about the demand at the start of each period, her only unknown demand parameter is λ_s .

Inspection rule: Depending on his valuation, each buyer on the market will choose to inspect the property (or not). An ‘inspection’ means that the buyer ‘visits the property and bargains with the seller’: once he visits he always bargains in the model. Before inspecting, the buyer observes a detailed advertisement about the property on the listing website and already knows his own valuation for it (v^b) without inspecting it.⁴ However, he only discovers the seller’s and potentially other inspecting buyers valuations if he meets and start to bargain with them when he inspects the house. Since the buyer suffers a cost of inspecting the property (I assume this cost to be infinitesimal for simplicity), he only does so if he expects to have a chance to buy it (and thus benefit from his inspection). With an infinitesimal inspection cost, it will be the case as long as $v^b > v^s$ (as there is always a chance for him to have no better competitor and to be able to buy the home in this case). Thus, to determine whether or not it is beneficial for him to enter, the buyer must build expectation about unknown v^s . The buyer has limited rationality and uses the list price p^L as a *signal* about v^s to build a naive conjecture that $\hat{v}^s = g(p^L)$.⁵ I use a simple affine functional form $g(x) = a_0 + a_1x$ with $0 < a_1 < 1$. It yields the following simple *inspection/entry rule* that all buyers follow:

$$\text{a buyer inspects the property if } v^b > g(p^L) = a_0 + a_1p^L$$

The seller only observes the number of inspections (N_t) and not the latent number of buyers on the market (N_t^{market}). Obviously, the seller can only sell her property to the inspecting buyers. She knows the buyers inspection rule $g()$, and since she knows that $v^b \sim \mathcal{LN}(\mu, \sigma)$, she is able to compute any

⁴An alternative story would be to say that visits are not costly at all (so all buyers visit and discover v^b when they do), but it is the decision to make a first offer which is costly. In this case the ‘inspection’ event would correspond to an ‘entry in bargaining’ or ‘offer to the seller’ instead of corresponding to a ‘visit’. Because of its cost (going to the bank and dealing with all financial details), buyers also need to decide whether or not to ‘inspect’ (make an offer). If one chooses to enter the bargaining, he learns about the seller and his competitors values. The final outcome is determined via the bargaining rule.

⁵Buyers are very naive and uninformed in this model. This can be justified by the fact that they are simple one-shot buyers, staying on the market only for one period. They do not have the time to gather information, thus they have a naive expectation.

Obviously, the seller’s list price choice is more complex than the buyers’ conjecture and depends only partially on v_s . In fact, the list price choice even depends on the buyers’ believed functional form $g()$ itself. I show in the results that this belief turns out not to be self-fulfilling: sellers reservation values differ from the buyers’ simple conjecture. Since the conjecture $g(p^L)$ might not be correct, i.e. $\hat{v}_s \neq v_s$, a buyer entry in the bargaining process does not necessarily result in a sale: in particular if $v^s > v^b > \hat{v}^s$, the buyer enters in a bargain with the seller, but both agents quickly realize that there will be no profitable trade for both of them, and no transaction occurs (as specified in the bargaining rule).

$Pr(v^b > g(p^L))$ for any $p^L > 0$. As a consequence, the homeowner faces a classic trade-off when she sets her list price; ceteris paribus, a high list price allows her to ‘sort’ buyers with a higher taste for her property, leading to a higher expected sale price. However, this also signals a higher reservation value to the buyers and thus reduces the probability that a buyer will visit and enter the bargaining process ($Pr(v_b > g(p^L))$), leading to a longer time on the market. As staying on the market is costly for the seller (she has to keep her home tidy, spend time for potential visits, etc.), the optimal list price, which maximizes the seller’s expected profit from sale, balances two opposite objectives: short time on the market and high sale price.

In addition to this classic trade-off, the list price also embeds an *informational externality*: ceteris paribus, a lower list price allows faster learning about λ_s . I detail the seller’s learning process in the next section.

3.3 Seller’s information and learning

When the seller decides to set her list price, she knows most parameters of the problem: she knows the bargaining rules, the buyers’ inspection rule, the distribution of buyers’ valuations (she does not know the realized value at the start of the period, but she knows that $v^b \sim \mathcal{LN}(\mu, \sigma)$ and knows μ and σ values), and that the number of buyers on the market will follow a $Poisson(\lambda_s)$ distribution. However, she has imperfect information about the demand since she *does not know the value of λ_s* : she only knows that in the population, $\lambda_s \sim Gamma(\alpha_0, \beta_0)$.

In period 0 (at the start of the listing), the seller forms an initial belief about λ_s , based on her rational expectation that $\lambda_s \sim Gamma(\alpha_0, \beta_0)$. Then for each period, she will update this belief via Bayesian learning rules, using the information she will acquire.

Let us drop the index s from λ_s and denote it only λ from now on (but remember it is an individual draw which is seller specific).

I determine the general *learning rule* for any period t . Suppose that the seller enters any period t with the prior belief that $\lambda \sim Gamma(\alpha_t, \beta_t)$ (i.e. $f_\lambda(\lambda) = \lambda^{\alpha_t-1} \frac{\beta_t^{\alpha_t} e^{-\beta_t \lambda}}{\Gamma(\alpha_t)}$, $\mathbb{E}[\lambda] = \alpha_t / \beta_t$ and $\mathbb{V}[\lambda] = \alpha_t / \beta_t^2$).⁶ To learn about λ , the seller will observe the number of inspections N_t (and not the latent number of buyers on the market N_t^{market} directly) in period t , and update her belief using this information.

First, recall that $N_t^{market} \sim Poisson(\lambda)$, and each of these buyers choose to inspect the property if their valuation is greater than $g(p^L)$. Thus, the process for the number of inspections depends on the list price of the period, $N_t \sim Poisson(\lambda Pr(v^b > g(p^L)))$: to learn about λ , the seller does not directly observe the latent process determined solely by λ , but by another process which is more or less close to

⁶In this model, λ is fixed and does not vary with time. However, making λ vary through time in an unknown way for the seller would not change the learning rule at all (the seller would still only be able to observe the number of inspections).

it depending on the choice of list price. Nonetheless, computing the posterior distribution remains quite simple in this case. For simplicity, we denote $Pr(v^b > g(p_t^L)) = c_t$. Then, the likelihood of observing $N_t = k$ inspections follows a $Poisson(\lambda c_t)$ distribution and is given by $f(N_t = k | \lambda c_t) = (\lambda c_t)^k e^{-\lambda c_t} / k!$. Observing N_t , the seller updates her initial belief using Bayes formula to compute the *posterior belief*:

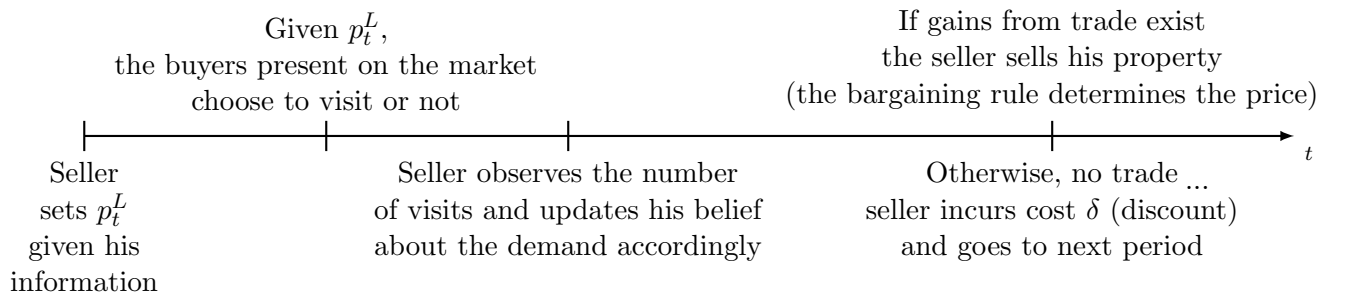
$$\begin{aligned}
f_\lambda(\lambda | N_t = k) &= \frac{f(N_t = k | \lambda c_t) f_\lambda(\lambda)}{f(N_t = k)} \\
&= \dots \\
&= \frac{(\beta_t + c_t)^{\alpha_t + k}}{\Gamma(\alpha_t + k)} \lambda^{\alpha_t + k - 1} e^{-(\beta_t + c_t)\lambda} \\
\iff \text{Posterior belief: } \lambda &\sim \text{Gamma}\left(\alpha_{t+1} = \alpha_t + N_t, \beta_{t+1} = \beta_t + Pr(v^b > g(p_t^L))\right)
\end{aligned}$$

The learning rule is fairly simple: the α parameter of the prior is updated by adding to it the observed number of inspections, while we add the individual probability of inspection to the β parameter.

In this context, the list price p_t^L embeds a *learning externality*. To see this, note that when the list price is so high that $Pr(v^b > g(p_t^L)) \rightarrow 0$, the seller will always observe $N_t = 0$, no matter what λ is. Her listing does not provide her any information about λ in this case, and the seller does not learn anything, her belief stays the same over the period (since $Pr(v^b > g(p_t^L)) = 0$, she will observe $N_t = 0 \forall N_t^{market}$ and thus $\alpha_{t+1} = \alpha_t$ and $\beta_{t+1} = \beta_t$). As p^L decreases, $Pr(v^b > g(p_t^L))$ increases and non-entries of some buyers are more and more likely due to the fact that there was indeed no buyer on the market (instead of being likely caused by a too low $Pr(v^b > g(p_t^L))$). This continues up to the opposite extreme scenario where p^L is so small that $Pr(v^b > g(p_t^L)) \rightarrow 1$, in which case a ‘non-entry’ only occurs when there is no buyer and $N_t = N_t^{market}$ (and the learning rule is actually equivalent to simple Bayesian updating with the basic $Poisson(\lambda)$ distribution where one adds 1 to β each period). In general, a smaller list price leads to a higher probability of inspection, which makes the ‘number of inspection process’ closer to the latent ‘number on market process’ and allows us to learn faster about the demand parameter (λ) determining the number of buyers on the market.

3.4 Seller’s optimization problem

Figure 3: Timeline of events in the model



The timing of the model is summarized in Figure 3. At the start of each period, the seller sets an optimal list price p_t^L in order to maximize her expected profit from the sale given her information set. Her information set consists of her belief about λ at the start of the period (which can be summarized by the two parameters (α_t, β_t)), and the knowledge of all the other parameters of the problem.

The number of buyers on the market is then drawn (from $Poisson(\lambda)$), as well as their valuations for the property. Given p_t^L and the inspection rule, each buyer chooses to inspect the property (or not). The seller then observes the number of inspections and updates her belief about λ according to the Bayesian learning rule: it determines her updated reservation value for the bargaining.

Once the seller has updated her v^s and all the buyers are entered, if at least one of the buyer has a valuation greater than the seller's, trade will occur according to the bargaining rule.⁷ Otherwise, there is no room for beneficial trade and no trade occurs, the seller incurs a cost of keeping her property on the market (δ under the form of a discount factor) and goes to the next period where she repeats the same process, starting with her updated belief.

I repeat this game over an infinite horizon up to the point where the seller sells her property.⁸

Denote $\Omega_t = (\alpha_t, \beta_t)$ the seller information set at time t . Also denote the seller value $v^s(\Omega_t, \mu, \sigma, a_0, a_1, \delta)$ simply as $v^s(\Omega_t)$. Notice that the seller's valuation does not depend on λ directly and only depends on what the seller believes λ to be. The true λ will only impact the updating process of this belief (by generating the true number of buyers on the market that the seller will observe).

This value is pinned down by the following *Bellman's equation* which represents the problem of the homeowner when she sets her list price optimally given her information set at the start of each period:

$$v^s(\Omega_t) = \max_{p_t^L} \sum_{k=0}^{\infty} \mathbb{E} \left[Pr(N_t = k | p_t^L, \lambda) \middle| \Omega_t \right] \mathbb{E} \left[\Pi^s(N_t = k, p_t^L, \Omega_t) \right]$$

where the expected probabilities of receiving k visits based on the starting beliefs (Ω_t) are

$$\mathbb{E} \left[Pr(N_t = k | p_t^L, \lambda) \middle| \Omega_t \right] = \int (\hat{x} P(v^b > g(p_t^L)))^k e^{-\hat{x} P(v^b > g(p_t^L))} / k! f_\lambda(\hat{x}) d\hat{x}$$

$$\text{with } f_\lambda(x | \Omega_t) = \frac{\beta_t^\alpha}{\Gamma(\alpha_t)} x^{\alpha_t - 1} e^{-\beta_t x}$$

and the corresponding profit function depends on the number of inspections and known updating in the

⁷The bargaining game is done within-period, meaning that future periods buyers can never enter it before the end of the process: it is as if I assumed that the sequential bargaining period was infinitely smaller than the dynamic game period of two weeks. The dynamic game only impacts the seller's reservation value that is built based on her expected gain (in the present or future): thus it only impacts her choice to leave the table in the bargaining process.

⁸In practice, to reduce the computational burden of the simulation, we choose a maximum number of weeks (e.g. two years) on the market, after which we stop the simulation if no sale has occurred.

case where this number of inspection indeed happens (i.e. $v^s(\Omega_{t+1})$ instead of Ω_t), as defined below:⁹

$$\Pi^s(N_t = k, p_t^L) = \begin{cases} \delta v^s(\Omega_{t+1}) & \text{if } v^s(\Omega_{t+1}) > v_{(k)}^b \\ \underbrace{p^s(v_{(k)}^b, v_{(k-1)}^b, v^s(\Omega_{t+1}), p_t^L)}_{\text{bargaining rule function}} & \text{otherwise} \end{cases}$$

with $\Omega_{t+1} = (\alpha_t + N_t, \beta_t + Pr(v^b > g(p_t^L)))$. With the special case that $\Pi^s(N_t = 0, p_t^L) = \delta v^s(\Omega_{t+1})$. The expectation of seller profit is taken with respect to the two highest buyers' values (which are the only ones which potentially matter in the bargaining rule, and which are unknown to the seller when she sets her list price) using the joint density of the two highest order statistics among k . This joint density of two order statistics is in general given for any $i < j \in 1, 2, \dots, n, \forall x < y \in \mathbb{R}$ by:

$$f_{(i,j):n}(x, y) = \frac{n!}{(i-1)!(j-i-1)!(n-j)!} [F(x)]^{i-1} [F(y) - F(x)]^{j-i-1} [1 - F(y)]^{n-j} f(x) f(y)$$

where f is the truncated lognormal(μ, σ) on $x > g(p^L)$, and F its cdf.

This value function is estimated via value function iteration. The iteration is done on a discrete grid of α and β values. Values for points inside the state space but out of the grids are approximated via bilinear interpolation between the four surrounding points.

4 Estimation method and identification

The structural parameters that I want to estimate are: $(\mu, \sigma, a_0, a_1, \delta, \alpha_0, \beta_0)$. Denote the structural parameters vector θ . θ is estimated via simulated method of moments (SMM). The idea of this estimation method is to find the set of parameters for which the simulated sellers' behavior will be the closest to the observed sellers' behavior. To do so, I select features from the empirical data that I want to reproduce by picking a vector of N empirical moments of interest. I denote m^d this $N \times 1$ vector of moments. Then, for a given θ , I construct the corresponding counterpart vector of simulated moments $m^{sim}(\theta)$. These simulated moments are computed on the selling outcomes data of S ($=100000$) iid simulations of my model with underlying structural parameters θ .¹⁰

I estimate the seven unknown parameters by minimizing a distance function between empirical and simulated moments such that:

$$\hat{\theta} = \underset{\theta}{argmin} \left[m^d - m^{sim}(\theta) \right]' W \left[m^d - m^{sim}(\theta) \right]$$

⁹For any number of visits that occur when the seller chooses to set a given p_t^L , she knows how she will update Ω_t to Ω_{t+1} , thus she knows the corresponding updated value that she would bargain with in each specific entry case ($v^s(\Omega_{t+1})$). The probability of observing each specific number of entries are computed using the period starting belief Ω_t .

¹⁰Recall that I normalized the data by the predicted sale price, which allows me to compare everything on the same scale by using a single representative problem for every seller (independent of the 'quality' of the property). Hence, I can run iid simulations of the representative problem and compare it to the normalized real data.

where W is a $N \times N$ positive definite weighting matrix equal to the inverse of the variance-covariance matrix of my moments (computed using bootstrap). To find $\hat{\theta}$ I use a controlled random search algorithm (Price, 1983).

To estimate the seven structural parameters I pick a list of $N = 61$ moments representing the features I want to reproduce with my model. These moments can be categorized in three main dimensions of the selling process: the time on the market, the distribution of sale price and the list price dynamics (see Appendix B for a list of all the moments). Most of the price moments are relative to the initial list price. This allows me to reduce issues caused by the ‘scaling’ of the problem or to soften the impact of unobserved heterogeneity not accounted for in the hedonic value estimation. I have only two ‘non-relative’ price moments: the average sale price and the average list price. They are compared to their counterparts normalized by the predicted financial value (estimated by hedonic regression) in the data.

These moments allow us to identify the parameters. Intuition about identification is non-trivial as most parameters influence several features of the model simultaneously. (α_0, β_0) determines the initial value and list price choice. By pinning down the λ distribution, they also determine the list price dynamics and have a strong influence on the time on the market. (μ, σ) pin down the buyer valuations, and thus the seller value and the list price level he can set. They also influence the final sale price through the buyer and the seller value. In particular, (μ, σ) will directly determine the distribution of sale price when it is above the final list price (as the model imply that in this case $p^s = v_{(n-1)}^b$ and these two parameters solely determine the distribution of $v_{(n-1)}^b$). a_0 and a_1 are also identified via the sale price distribution. In particular, they determine the minimum level of entry of buyers and thus the minimum ratio of sale over list prices. δ only impacts the seller valuation and allows to adjust it better than $(a_0, a_1, \mu$ and $\sigma)$ which face more restrictions (as they determine more specific moments).

5 Results

5.1 Value function and optimal list price

Figure 4: $v^s(\alpha, \beta)$

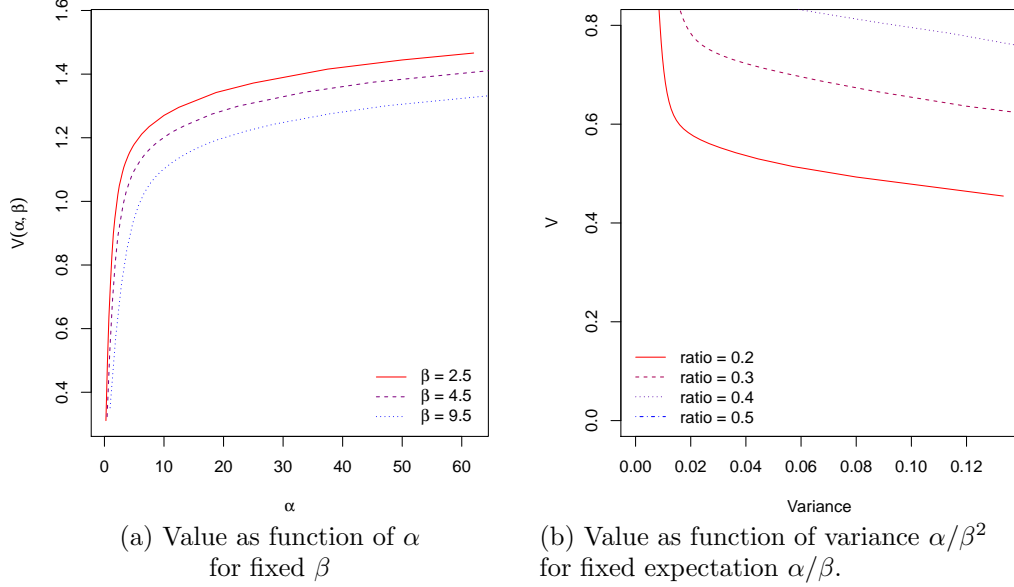
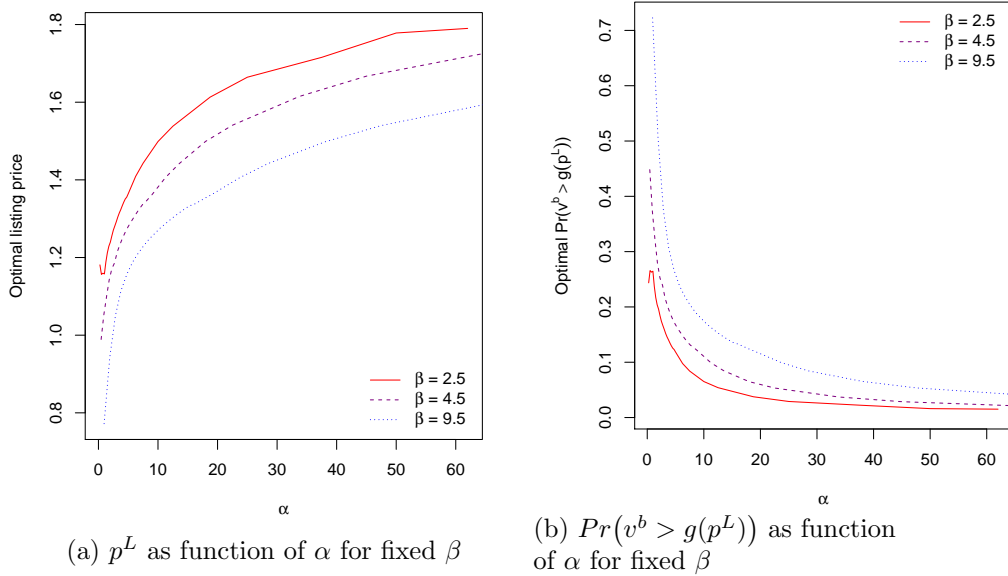


Figure 5: Optimal list price $p^L(\alpha, \beta)$ and corresponding $Pr(v^b > g(p^L))$



I obtain intuitive results for the value function (Figure 4) and optimal list price (Figure 5).

For the interpretations, recall that if $\lambda \sim \text{Gamma}(\alpha, \beta)$, then $\mathbb{E}(\lambda) = \alpha/\beta$ and $V(\lambda) = \alpha/\beta^2$.

First, the value increases with the expected number of buyers on the market. Second, at the fixed average belief, the less uncertain the seller is (i.e., the smaller variance of the belief), the higher the value

she obtains: as the uncertainty decreases, the value converges to the perfect information benchmark (as if the seller knew its own λ draw). As for the optimal list price, the higher the expected latent number of buyers on the market (higher expected λ), the smaller the chosen probability of inspection (via a higher list price), which balances the expected number of inspections overall (depending on $\lambda Pr(v^b > g(p^L))$).

5.2 Parameters and Moments

Table 2: Parameter Estimates of the Structural Model

Parameter	Description	Estimate	Std. Errors
δ	Subjective discount factor (and cost of listing)	0.986	0.00009
<u>Demand parameters</u>			
<i>Valuation process:</i>			
	$v^b \sim \mathcal{LN}(\mu, \sigma)$		
μ	Mean of buyer valuation	-0.0425	0.00044
σ	Standard deviation of buyer valuation	0.1902	0.00039
<i>Number of buyers:</i>			
<i>(and initial rational belief)</i>			
	$N^{market} \sim Poisson(\lambda)$ and $\lambda \sim Gamma(\alpha_0, \beta_0)$		
α_0	α prior belief and true distribution	4.48	0.005
β_0	β prior belief and true distribution	11.21	0.0057
<u>Inspection rule</u>			
	buyer inspects if $v^b > a_0 + a_1 p^L$		
a_0	Buyer's conjecture about seller reservation value: constant	0.409	0.0018
a_1	Buyer's conjecture about seller reservation value: slope	0.58	0.0015

Table 2 reports the parameter estimates. To understand the mechanism at play, we focus on the screening of buyers implicitly done by the seller at the optimal parameters. Take the screening in the initial period for example. The seller face a demand $v^b \sim \mathcal{LN}(\mu = -0.0425, \sigma = 0.19002)$. She sets an initial listing price of 1.106 (no heterogeneity in the first period since it is a representative agent model and everyone starts with the same rational expectation). In this case, $g(p^L) = 1.05021$, $Pr(v^b > g(p^L)) = 31.52\%$: the seller aims for high quantiles of the demand (top 31.52%). For example, if λ is drawn at its average value ($\alpha_0/\beta_0 = 0.4$), then the true rate of inspection $\lambda Pr(v^b > g(p^L))$ is equal to 12.61%.

The sellers with a high draw of λ will have a considerably higher probability of sale than they expect, and will thus spend a shorter time on the market than they would like to if they had perfect information (however, some will stay on the market and increase their price in the next period if they received multiple visits, they often choose to sell at a price they would have refused with perfect information

in the end). On the contrary, sellers with a low draw of λ have a very small probability of sale and are aiming at too high quantiles of the demand (compared to what they would do if they had perfect information about λ): they will ‘survive’ longer on the market (than people with high λ). As they stay on the market and obtain information that their λ is low, they will decrease their list price to aim at lower quantiles of the buyer valuations in order to compensate for their lower λ and still have a chance to sell their property.

The model is able to match the decreasing list price dynamics observed in the data thanks to the fact that sellers who stay longer on the market (‘survivors’) are sellers with low draw of λ . These sellers progressively learn and adjust their expectation about the demand (λ) downwards, and thus decrease their list price choice. For these individuals with low draws of λ , the initial belief was too ‘optimistic’: the correction of their initial belief via learning is the reason for decreasing list price.


In addition to this, the decreasing list price is also due to some selection/survivor effect (even for high draws of λ): those who are unlucky (no matter what λ is) and observe no entries will stay longer on market and have a decreasing belief. Thus, in general, those who stay longer on the market are likely to have observed fewer entries (even though some stay because they refused offers, the majority stays because they received none), and thus have a more ‘downward updated’ belief.

Table 3: Actual and Simulated Moments


Moment	Actual	Simulated
Mean sale price	1.008	1.02
Mean ratio sale/final listing price	0.955	0.962
Mean initial list price	1.107	1.106
% of accepted offers equal to list price	0.15	0.235
% of accepted offers below list price	0.734	0.712
Mean week on the market (knowing that <52 weeks)	14.817	14.736

Table 3 and Figure 6 illustrate how the model matches the moments. Overall, for a small number of parameters (7), it fits the data reasonably well. The list price dynamics and the distribution of the time on the market are well fitted. The sale price distribution (relative to the list price) is matched correctly, except for the tails. In particular, I fail to reproduce the number of sales above the list price (only 5.3% in the simulation against 11.6% in the data). This is because, even with the split demand, the model is still unable to match the time on the market distribution and the sales above the list price at the same time.


Figure 6: Actual and Simulated Moments



(a) Proportion of unsold properties as a function of time on the market



(b) Distribution of sale prices (ratio to final list price)



(c) Average list price as a function of time on the market (ratio to initial list price)

6 Value of information

All of the comparisons in this section are completed using the model at the estimated optimal parameters given previously.

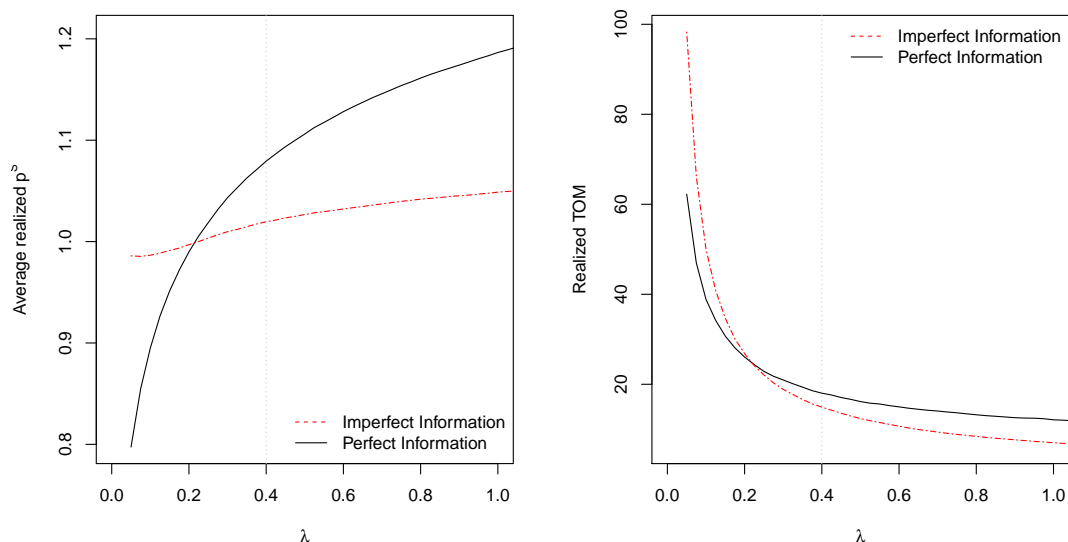
I use my estimated model to get an idea of the *value of information*: how much the seller would gain from being better informed? To answer this, I compare the imperfect information outcomes to the outcomes obtained by a *perfectly informed* (denoted PI) seller, who would know the value of her λ draw.

I compare this benchmark to the value obtained by an imperfectly informed (denoted II) seller who does

not know her draw λ and starts from rational expectation $\text{Gamma}(\alpha_0, \beta_0)$. To compute the ‘realized value’ of this imperfectly informed seller, I simulate the complete selling process (with corresponding dynamic belief updating) and observe the price p^S and time spent on the market by the seller. From period 0, the average (over several simulations) discounted sale price ($\delta^t p^S$) gives the empirical counterpart to the theoretical seller value.

To establish an idea of the value of information, I compare the perfectly informed benchmark valuations to the counterparts’ realized outcomes of imperfectly informed sellers. In particular, I do this for several values of λ . In this way, I can see how far II sellers end up from the perfect information benchmark as a function of their initial information error (how far the initial rational belief was from their true λ).

Figure 7: Value of Information: Sale Price and TOM
(average computed over 10000 simulations)



(a) Average obtained p^S (not discounted) as a function of λ (for fixed initial belief (α_0, β_0)) (b) Average time spent on market (TOM) as a function of λ (for fixed initial belief (α_0, β_0))

Figure 7a shows the average sale price and time on the market (simulated) for different values of λ . This is done for a PI seller who knows λ , and for a II seller who starts from initial rational belief that $\lambda \sim \text{Gamma}(\alpha_0, \beta_0)$ (calibrated at the model estimates) and who progressively updates it.

First for both types of sellers, I observe that the higher the λ the higher the sale price obtained and the shorter the time spent on market.

More interestingly, we see that the sale price reacts more to λ for perfectly informed sellers. Indeed, the average sale price obtained by a PI seller varies from 0.797 when $\lambda = 0.05$, to 1.187 when $\lambda = 1$. While the average sale price obtained by a II seller (with initial belief $\text{Gamma}(\alpha_0, \beta_0)$) varies from 0.986 when $\lambda = 0.05$, to 1.049 when $\lambda = 1$. The fact that imperfectly informed agents start from the

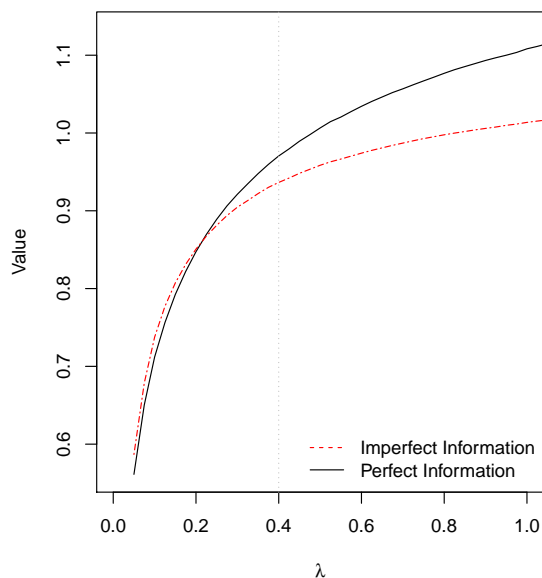
same initial belief smooths the obtained sale price: at the time they sell, despite their updating, they are generally still far from perfect information. Thus, their list price is far from the optimal one (if perfectly informed they would set a lower price to sell faster), hence the final sale price difference. In terms of sale price, sellers with bad draws of λ benefits from this, while sellers with high draws of λ get a smaller price than what they could.

However, as shown in Figure 7b, imperfectly informed sellers with bad draws are able to earn more only because they spend too much time on the market (about 98 weeks on average for a II seller with draw $\lambda = 0.05$, compared to 62 weeks for her PI counterpart). In terms of value obtained by the II seller, this overly-long time on the market attenuates the higher sale price obtained.

Similarly, imperfectly informed sellers with high draws spend a shorter time on the market than their perfectly informed counterparts, since they put their listing at a suboptimal list price (and screen less the demand than they would if they were perfectly informed). This could offset the loss in the sale price they endure compared to the PI benchmark.

The question now is to determine how the effects on sale price and time on the market translates into the average ‘value’ (average discounted sale price) obtained (from period 0).

Figure 8: Value of Information: Information Paradox
(average computed over 10000 simulations)



Value as a function of λ (for fixed initial belief (α_0, β_0)).
The vertical grey line represents $\mathbb{E}[\lambda]$ at the initial rational belief (i.e. α_0/β_0).

Figure 8 shows how these two effects translate in terms of value. Obviously, the higher the demand λ , the higher the value obtained for perfectly and imperfectly informed agents. However, one can notice

an *informational paradox*.

Indeed, one would expect the perfectly informed agent to always outperform the imperfectly informed counterpart because she is solving the ‘true’ problem (knowing λ), and thus optimizes correctly. However, this is not always the case, and the imperfectly informed agent can perform better than her perfectly informed counterpart in this model. Indeed, there is an area where the value obtained by the imperfectly informed agent who starts from the belief that $\lambda \sim \text{Gamma}(\alpha_0, \beta_0)$ is higher than that obtained by her perfectly informed counterpart who knows λ (i.e. the area where the red-dotted curve is above the black curve in Figure 8). This happens only to agents with bad draws (i.e., imperfectly informed agents who expect λ to be considerably higher than what it truly is). The realized outcomes (discounted sale price) with an ‘overoptimistic’ belief can be higher than the ones the agent could obtain if she were perfectly informed and optimizing (choosing her list price) knowing the exact value of λ . This means that having better information about λ is not necessarily beneficial for the seller, as the value of additional information may indeed be negative.

This seems somewhat perplexing, as one should not be able to do better than a perfectly optimizing agent who knows exactly what demand she should expect, and thus, what her true value v^s should be. The explanation lies in the estimation of her reservation value by the seller. Indeed, even if the overoptimistic seller is not optimizing correctly with respect to the true λ (she sets a too high p^L and stays too long on the market by ‘screening’ too much and ‘over-rejecting’ some buyers offers), she genuinely *overestimates her reservation value* $v^s(\alpha, \beta)$. This gives her a better ‘bargaining position’ (the threshold at which she leaves the bargaining game) in the bargaining game, which allows her to obtain a higher sale price than the one a PI seller would obtain if she was trading with the same buyers. One can directly see this in the bargaining rule: ceteris paribus, if v^s is higher, $v^s + 0.5(v_{(n)}^b - v^s) = 0.5v^s + 0.5v_{(n)}^b$ is also higher. There is some level of overconfidence where the overconfident seller obtains a sale price sufficiently high to offset the longer time she spends on the market, yielding a better outcome overall. At some point (not visible in Figure 8), the gains in the bargaining position are offset by a too large ‘expectation error’. By being ‘too overoptimistic’ the agents spend too long time on the market, which offset the stronger bargaining position and yields a lower value than being perfectly informed.¹¹

Notice that the highest values in case of overoptimism are only explained by this improved bargaining position. To check this, I can recompute a variant of the model where the seller is infinitely more patient than the buyers in the bargaining game. This way she has all the bargaining power and can

¹¹It does not appear with our estimated set of parameters, but with other parameters, for very small λ draws, the black curve will again be higher than the red one. There is an ‘optimal level’ of ‘overconfidence’ (not too large or too small) for which the misinformed agent performs better than her perfectly informed counterpart. You can see this by zooming in on Figure 8: the positive gap between II and PI values (when λ is low) has an inverted U shape.

always extract the highest buyer valuation ($v_{(n)}^b$) without splitting the pie. Thus, she already has full bargaining power and cannot have a stronger bargaining position due to her overconfidence. In this case, the information paradox disappears: the perfectly informed agent always performs better (not displayed here).¹²

This explains that while the imperfectly informed with a high draw of λ is, on the other hand, considerably worse off than her perfectly informed counterpart, she does not only commit optimization errors (setting a too low list price, resulting in less screening of buyers, lower sale price and shorter time spent on market) due to her imperfect information set, but she also has a weaker bargaining position because she is too ‘pessimistic’ in her reservation value estimation v^s .

Thus, the cost of imperfect information is higher for people facing high demand. High demand sellers would prefer to be perfectly informed of the demand that they face, while on some levels, the low demand sellers can be better off by being overoptimistic and ignorant that they face a low demand.

7 Conclusion

Taking advantage of a new large dataset of real estate listings, I highlight some evidence that imperfect information and sellers’ learning impact the selling outcomes on the housing market. I have developed a simple theoretical model with a new Bayesian learning application in order to explain some of the housing market stylized facts.

Learning about the demand by rational but imperfectly informed home-sellers is a key feature of the model used to explain these facts. In particular, the progressive downward adjustment/correction of the belief of individuals facing low demand is key to explaining the decreasing list price dynamics, and matches well the distribution of time spent on market. My work also highlights a paradox that the value of information is not necessarily positive. Indeed, by being imperfectly informed and overconfident about the demand, a seller can overestimate her reservation value and have a stronger bargaining position. This allows her to extract a higher sale price, which can compensate her mis-optimization and longer time spent on the market (with respect to her perfectly informed counterpart).

This theoretical work could serve as the foundation for future applications relative to the home-selling problem. For example, it can be easily extended to study learning within complete neighborhoods, or to study dynamic entry/withdrawal decisions of listings by the sellers to match market stocks.

¹²Technically, when the seller is infinitely more patient, her reservation value (v^s) vanishes from the bargaining equation determining the sale price. For a given buyer value (v^b), v^s now only impacts the decision to reject an offer, without impacting the sale price. Thus, it will only lead to a misguided rejection decision and the stronger bargaining position will not matter (because the seller already has all the power). To be clear, the seller will still get a higher average sale price (because she screens higher value buyers with her rejections). However, the longer time on the market will offset this advantage and the PI seller will always be better off.

reject the offer, they go to the next period (with delay τ) where they make simultaneous counteroffers, as in the first period, and the game repeats itself, etc.

A.2 Solution

This game always has a unique subgame-perfect equilibrium outcome where the good is sold immediately to buyer H (the seller accepts his offer directly) at a price:

- $p = \frac{\delta}{1+\delta}$ if $\frac{\delta}{1+\delta} \geq v$. Notice that this is the bilateral bargaining price of the seller with buyer H, in this case, it is as if the second buyer was absent (he represents a non-credible threat for buyer H / non-credible outside option for the seller).
- $p = v$ if $\frac{\delta}{1+\delta} < v$. In this case the presence of the buyer L matters, gives more ‘power’ to the seller, and forces buyer H to pay a higher price than if buyer L was absent.

A.3 Proof

As in the classic [Shaked and Sutton \(1984\)](#) proof, let m_b and M_b be the infimum and supremum of equilibrium payoffs to the buyer H in the game. Let m_s and M_s be the infimum and supremum of equilibrium payoffs to the seller in the companion game in which she would move first (i.e. starting from period $t = 1$, for example). Let us also assume for now that the buyer L always makes the same equilibrium offer denoted s (≤ 1 since $v_L \leq 1 = v_H$). From the point of view of bilateral bargaining between the seller and buyer H, it acts as an ‘outside option’ for the seller: if she accepts it, she obtains s and she ‘leaves’ buyer H with nothing.¹⁵

As in the bilateral bargaining with outside option proof from [Binmore et al. \(1989\)](#), I have the following system of inequalities which hold:

$$m_b \geq v - \max\{\delta M_s, s\} \quad (1)$$

$$v - M_b \geq \max\{\delta m_s, s\} \quad (2)$$

$$m_s \geq v - \delta M_b \quad (3)$$

$$v - M_s \geq \delta m_b \quad (4)$$

Inequality (1) can be explained as follows: the seller must accept any opening offer greater than what she can get by making a counteroffer to buyer H in the next period, or by accepting the low buyer offer (thus no δ cost). As a consequence, the buyer H cannot get less than $v - \max\{\delta M_s^H, s\}$, hence the first inequality. Inequality (2) follows from the fact that the seller must get at least either δm_s^H by

¹⁵Obviously since buyer H can give more than L (since $v_H = 1 \geq v_L$), he is always able to attract the seller to trade with himself, by offering something greater or equal than s if necessary.

making a counteroffer to buyer H, or s by accepting the low buyer offer. As a consequence, the buyer H can get at most $M_b \leq v - \max\{\delta m_s^H, s\}$, hence the second inequality. Similarly, inequality (3) and (4) comes from the same reasoning but for buyer H and thus, there is no s involved (as if $s = 0$, i.e. no offer from another seller that he could accept for example/no ‘outside option’).

Now, to determine the equilibrium outcomes, distinguish three cases:

- If $s \leq \delta m_s$: (\iff the offer from L is irrelevant)

Combining (1) and (4) yields:

$$\delta - \delta^2 M_s \leq \delta m_b \leq 1 - M_s, \text{ thus } \delta - \delta^2 M_s \leq 1 - M_s, \text{ which gives: } M_s \leq 1/(1 + \delta).$$

Combining (2) and (3) (rewritten) yields:

$$1 - m_s \leq \delta M_b \leq \delta - \delta^2 m_s, \text{ thus } 1 - m_s \leq \delta - \delta^2 m_s, \text{ which gives: } 1/(1 + \delta) \leq m_s.$$

Thus:

$$\frac{1}{1 + \delta} \leq m_s \leq M_s \leq \frac{1}{1 + \delta}$$

Similarly, combining (2) and (3) for the upper bound, and (1) and (4) for the lower bound, yields:

$$\frac{1}{1 + \delta} \leq m_b \leq M_b \leq \frac{1}{1 + \delta}$$

Thus, $m_s = M_s = m_b = M_b = 1/(1 + \delta)$ in this case (and buyer should offer $1/(1 + \delta)$ to the seller, who will accept in this case). Thus, this case should happen when $\delta m_s = \delta/(1 + \delta) \geq s \iff s \leq \delta/(1 + \delta)$.¹⁶

- If $\delta m_s < s < \delta M_s$: (2) becomes: $1 - M_b \geq s > \delta m_s$, but we still have, as before: $1 - M_b \geq \delta m_s$, thus we will still find $\frac{1}{1 + \delta} \leq m_s \leq M_s \leq \frac{1}{1 + \delta}$, which is a contradiction.

Thus, this case is not possible.

- If $\delta M_s \leq s$: (\iff the offer from L is greater than what the seller could get with classical bilateral bargaining with H)

For the buyer, we directly have from (2) that: $M_b \leq 1 - s$, and from (1) $m_b \geq 1 - s$. Thus:

$$1 - s \leq m_b \leq M_b \leq 1 - s \iff m_b = M_b = 1 - s$$

¹⁶Intuitively, instead of resorting to inequalities, see the proof as in [Shaked and Sutton \(1984\)](#) in each case. Let’s take an example with the proof for the upper bound of M_b .

We focus on the subgame starting from period $t = 2$. The game which starts at this point is the same as the initial game (its first repetition) but with a discounted sum of payoffs $= \delta^2$ (cannot get more than this). By definition, the buyer can get at most $\delta^2 M_b$ at this point. Now, consider the (companion) subgame starting in the preceding period $t = 1$. Any offer by the seller which gives the buyer more than the supremum of its payoffs ($\delta^2 M_b$) should be accepted. So there is no perfect equilibrium in which the buyer receives more than $\delta^2 M_b$, and thus it follows that the seller should get at least $\delta - \delta^2 M_b$ in this period (it is $\delta - \delta^2 M_b$ and not $1 - \delta^2 M_b$ since the discounted value of the total payoff at time $t = 1$ is δ and not 1). In other words, $m_s \geq \delta - \delta^2 M_b$. As a consequence, starting in period $t = 0$, the seller will not accept anything less than the infimum of what she will receive in the game beginning next period (which has present value $\delta - \delta^2 M_b$). Thus, the buyer can get, at most $M_b \geq 1 - \delta + \delta^2 M_b$. This finally gives the upper bound: $M_b \leq 1/(1 + \delta)$.

Obtain that $m_b \geq 1/(1 + \delta)$ and the results for the seller by similar reasoning.

Which means that the buyer should immediately make an offer of s to the seller, and he will not be able to obtain more.

We still need to compute the seller outcomes (to check when this case happen).

As before, from (1) and (4) we have: $M_s \leq 1 - \delta(1 - s)$. From (2) and (3) we have: $1 - \delta(1 - s) \leq m_s$.

$$1 - \delta(1 - s) \leq m_s \leq M_s \leq 1 - \delta(1 - s) \iff m_s = M_s = 1 - \delta(1 - s)$$

Thus, this case should happen when $\delta M_s = 1 - \delta(1 - s) \leq s \iff s \leq \delta/(1 + \delta)$.

Therefore, if subgame-perfect equilibria exist, they generate a unique subgame-perfect equilibrium outcome. In addition, existence is trivial: each player always demands his equilibrium payoff when proposing, and accept his equilibrium payoff (or more) when responding.

Thus, we have the result that the buyer H makes an offer at p , which is immediately accepted by the seller with the offer from the low buyer s as her outside option. With p defined as:

$$p = \begin{cases} 1 - 1/(1 + \delta) = \delta/(1 + \delta), & \text{if } s \leq \delta/(1 + \delta) \\ s, & \text{if } s \geq \delta/(1 + \delta) \end{cases}$$

The result is quite intuitive: either the offer from buyer L is too low and is not taken into account by the buyer H and the seller (irrelevant outside option for the seller, they do classic bilateral bargaining), or it is high enough and allows the seller to gain a credible threat, which increases her payoff.

Now the question is to determine what is s , the equilibrium offer from the buyer L (if it exists).

Let's assume that the buyer L cannot make an offer greater than his valuation v .¹⁷

¹⁷Even if it seems natural, this is an important assumption to get rid of absurd equilibria where L would bid more than v in this setup. Because the buyers bids simultaneously with perfect information, it is close to the case of first price sealed bid auction with perfect information without restrictions on bids. It is well known that the set of Nash equilibria of this kind of auction is determined by three conditions: it is the set of profiles b of bids with $b_H \in [v_L, v_H] = [v, 1]$, $b_j \leq b_H \forall j \neq H$ and $b_j = b_H$ for some $j \neq 1$. Thus we could have any equilibrium offer $> v$ from the low buyer, and the proof may fall down.

On the other hand, if we impose that he cannot bid more than his valuation v - which makes sense in the context of bargaining where the set of possible offers is between the valuation of the seller and the one of the corresponding buyer - then the only Nash equilibrium is $b_H = b_L = v$ (and because of the tie breaking rule, H wins). Indeed, it is clear that any outcome with $b_H < v$ and $b_L < v$ is not an equilibrium since one of the two players would gain more by increasing his bid. Similarly: $b_H = v$ and $b_L < v$ is not an equilibrium either, since in this case H would be tempted to decrease slightly his offer in order to get more.

Otherwise, instead of making this assumption, one solution would simply be to make the buyers bid non simultaneously: H first, then L, in which case L would never have interest to bid more than H's offer if it's higher than v , and thus H would never bid more than v in a first price auction... One has to choose between the sequential bids or the simultaneous bids with the natural assumption that the buyers cannot bid more than their value. I prefer the former since it seems unnatural to make them bid sequentially with predefined order based on the valuations.

- If $\delta/(1 + \delta) > v$ then anyway L has no power to disturb bilateral bargaining between the seller and H, s does not matter at all in the problem (irrelevant outside option). Buyer L could make any offer $s \leq v$ in equilibrium, it would not change the equilibrium outcomes, let's assume he bids $s = v$ in this case.
- If $\delta/(1 + \delta) < v$ then L has some power (it is close to first price sealed bid auction with perfect information in this case).

We find that $s = v$ in this case is the only equilibrium. Indeed: if $v > s$, it is possible that buyer H wins the auction by bidding b_H below v and above s which is not an equilibrium since in this case L would be better off by increasing his bid above b_H . At the same time, it is not possible that buyer L wins the auction in equilibrium (buyer H can always bid more). The only equilibrium offer from L is $s = v$.

Thus, as expected, in equilibrium the low buyer will bid his valuation $s = v$. This yields the final result that the only equilibrium payoffs is that buyer H makes an offer at p which is immediately accepted by the seller (who have an offer v from the low buyer as outside option). With p defined as:

$$p = \begin{cases} \delta/(1 + \delta), & \text{if } v \leq \delta/(1 + \delta) \\ v, & \text{if } v \geq \delta/(1 + \delta) \end{cases}$$

A.4 Extension to N buyers

One can easily generalize to more than two buyers. Indeed, we do not really care about additional buyers choices: as long as they are lower than v , they will not have any impact on the equilibrium outcomes. Only the two highest valuations matter, so simply consider that in the proof here; that buyer H and L are the two buyers with the highest valuations and the proof is already generalized for any number of buyers.

A.5 Extension: list price

I slightly modify the problem by adding an exogenous list price p^l . The list price serves as a commitment device; thus, if a buyer makes an offer greater or equal to the list price in the first period, the seller is obliged to accept to sell her good. If two buyers make offers greater than the list price, she obviously accepts the greatest offer (and in case of tie, she chooses buyer H). The list price only works in the first period; after which, the game is unchanged (in period $t = 2$ we go back to previous case where the seller can refuse an offer greater than p^L).

rescale the problem to correspond to the one in the article, we get the same bargaining rule.

Appendix B Complete list of actual and simulated moments

Table 4: Actual and Simulated moments (complete table)

Moment	Actual	Simulated
Mean sale price	1.008	1.02
Mean ratio sale/final listing price	0.955	0.962
Mean initial list price	1.107	1.106
% of accepted offers equal to list price	0.15	0.235
% of accepted offers below list price	0.734	0.712
Mean week on the market (knowing that <52 weeks)	14.817	14.736
% unsold after 2 weeks	0.911	0.88
% unsold after 4 weeks	0.802	0.78
% unsold after 6 weeks	0.709	0.692
% unsold after 8 weeks	0.626	0.608
% unsold after 10 weeks	0.546	0.534
% unsold after 12 weeks	0.482	0.474
% unsold after 14 weeks	0.425	0.415
% unsold after 16 weeks	0.376	0.366
% unsold after 18 weeks	0.332	0.323
% unsold after 20 weeks	0.294	0.284
% unsold after 22 weeks	0.26	0.254
% unsold after 24 weeks	0.23	0.227
% unsold after 26 weeks	0.202	0.201
% unsold after 28 weeks	0.174	0.178
% unsold after 30 weeks	0.153	0.158
% unsold after 32 weeks	0.135	0.14
% unsold after 34 weeks	0.12	0.124
% unsold after 36 weeks	0.105	0.11
% unsold after 38 weeks	0.093	0.097
% unsold after 40 weeks	0.082	0.087
% unsold after 42 weeks	0.073	0.079
% unsold after 44 weeks	0.064	0.07
% unsold after 46 weeks	0.056	0.062
% unsold after 48 weeks	0.049	0.056
% unsold after 50 weeks	0.043	0.051
Mean list price in week 3	0.997	0.996
Mean list price in week 5	0.99	0.989
Mean list price in week 7	0.982	0.981
Mean list price in week 9	0.973	0.974
Mean list price in week 11	0.966	0.967
Mean list price in week 13	0.958	0.96
Mean list price in week 15	0.952	0.954
Mean list price in week 17	0.946	0.949
Mean list price in week 19	0.94	0.944

Mean list price in week 21	0.935	0.932
Mean list price in week 23	0.931	0.922
Mean list price in week 25	0.927	0.923
Mean list price in week 27	0.922	0.915
Mean list price in week 29	0.917	0.906
% sales/listing <0.70	0.014	0
% sales/listing <0.80	0.036	0
% sales/listing <0.85	0.064	0
% sales/listing <0.90	0.129	0.053
% sales/listing <0.92	0.186	0.213
% sales/listing <0.94	0.279	0.368
% sales/listing <0.95	0.344	0.436
% sales/listing <0.96	0.426	0.505
% sales/listing <0.97	0.522	0.564
% sales/listing <0.98	0.618	0.619
% sales/listing <0.99	0.697	0.669
% sales/listing <1.00	0.886	0.947
% sales/listing <1.02	0.949	0.959
% sales/listing <1.05	0.981	0.973
% sales/listing <1.10	0.992	0.988
% sales/listing <1.20	0.997	0.997

References

- ANENBERG, E. (2016): “Information frictions and housing market dynamics,” *International Economic Review*, 57.
- ARNOLD, M. (1999): “Search, Bargaining and Optimal Asking Prices,” *Real Estate Economics*, 27, 453–481.
- BINMORE, K., M. J. OSBORNE, AND A. RUBINSTEIN (1992): “Chapter 7 Noncooperative models of bargaining,” Elsevier, vol. 1 of *Handbook of Game Theory with Economic Applications*, 179 – 225.
- BINMORE, K., A. SHAKED, AND J. SUTTON (1989): “An Outside Option Experiment,” *The Quarterly Journal of Economics*, 104, 753–770.
- CARRILLO, P. E. (2012): “An Empirical Stationary Equilibrium Search Model Of The Housing Market,” *International Economic Review*, 53, 203–234.
- (2013): “To Sell or Not to Sell: Measuring the Heat of the Housing Market,” *Real Estate Economics*, 41, 310–346.
- CASE, K. E. AND R. J. SHILLER (1988): “The behavior of home buyers in boom and post-boom markets,” *New England Economic Review*, 29–46.
- (2003): “Is There a Bubble in the Housing Market?” *Brookings Papers on Economic Activity*, 34, 299–362.
- CHEN, Y. AND R. W. ROSENTHAL (1996a): “Asking Prices as Commitment Devices,” *International Economic Review*, 37, 129–155.
- (1996b): “On the Use of Ceiling-Price Commitments by Monopolists,” *RAND Journal of Economics*, 27, 207–220.
- DE WIT, E. R. AND B. VAN DER KLAUW (2013): “Asymmetric information and list-price reductions in the housing market,” *Regional Science and Urban Economics*, 43, 507–520.
- HAN, L. AND W. C. STRANGE (2014): “Bidding Wars for Houses,” *Real Estate Economics*, 42, 1–32.
- (2015): *The Microstructure of Housing Markets*, Elsevier, vol. 5 of *Handbook of Regional and Urban Economics*, chap. 0, 813–886.

- (2016): “What is the role of the asking price for a house?” *Journal of Urban Economics*, 93, 115–130.
- HAURIN, D. (1988): “The Duration of Marketing Time of Residential Housing,” *Real Estate Economics*, 16, 396–410.
- HOROWITZ, J. (1992): “The Role of the List Price in Housing Markets: Theory and an Econometric Model,” *Journal of Applied Econometrics*, 7, 115–29.
- KNIGHT, J. R. (2002): “list price, Time on Market, and Ultimate Selling Price: Causes and Effects of list price Changes,” *Real Estate Economics*, 30, 213–237.
- LIU, C. H., A. NOWAK, AND S. ROSENTHAL (2014): “Bubbles, Post-Crash Dynamics, and the Housing Market,” Working Papers 14-18, Department of Economics, West Virginia University.
- MERLO, A. AND F. ORTALO-MAGNE (2004): “Bargaining over residential real estate: evidence from England,” *Journal of Urban Economics*, 56, 192–216.
- MERLO, A., F. ORTALO-MAGNE, AND J. RUST (2015): “The Home Selling Problem: Theory And Evidence,” *International Economic Review*, 56, 457–484.
- MILLER, N. AND M. SKLARZ (1987): “Pricing Strategies and Residential Property Selling Prices,” *Journal of Real Estate Research*, 2, 31–40.
- ODEAN, T. (1998): “Volume, Volatility, Price, and Profit When All Traders Are Above Average,” *The Journal of Finance*, 53, 1887–1934.
- PIAZZESI, M. AND M. SCHNEIDER (2009): “Momentum Traders in the Housing Market: Survey Evidence and a Search Model,” *American Economic Review*, 99, 406–411.
- PRICE, W. L. (1983): “Global optimization by controlled random search,” *Journal of Optimization Theory and Applications*, 40, 333–348.
- RUBINSTEIN, A. (1982): “Perfect Equilibrium in a Bargaining Model,” *Econometrica*, 50, 97–109.
- SALANT, S. (1991): “For Sale by Owner: When to Use a Broker and How to Price the House,” *The Journal of Real Estate Finance and Economics*, 4, 157–73.
- SHAKED, A. AND J. SUTTON (1984): “Involuntary Unemployment as a Perfect Equilibrium in a Bargaining Model,” *Econometrica*, 52, 1351–64.