Imperfect Information, Learning and Housing Market Dynamics*

Christophe BRUNEEL[†]

Toulouse School of Economics

January 19, 2018

Abstract

This paper examines the problem of a homeowner who wants to maximize her expected profit from the sale of her property under uncertainty about the market conditions. Using a large dataset of real estate transactions in Pennsylvania between 2011 and 2014, I provide evidence that information frictions influence the selling outcomes on the housing market. I develop a theoretical dynamic search model of the home selling problem with imperfect information concerning the demand and Bayesian learning. The estimated model suggests that a high initial overestimation of the demand by the sellers, progressively adjusted via learning about the true market conditions, explains key features of the housing data such as the observed decreasing list price dynamics and the time on market. By comparison with a perfect information benchmark, my model exhibits an unexpected feature: the value of information is not always positive because an overconfident seller can obtain 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.

*Thanks to my advisor Thierry Magnac for his regular feedback. I also thank John Rust, Michel Le Breton and the participants in TSE econometrics and applied microeconomics workshops for their comments.

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

1 Introduction

Real estate transactions involve large financial amounts (\$206k on average in Pennsylvania) and typically 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. Yet making a succesful sale is tricky. Contrary to many other markets, the homeowner does not post a non negotiable price directly. Instead she chooses a list price in order to attract buyers and she commits to it (she must accept any offer above it) while bargaining with them to determine the final sale price afterwards (Horowitz, 1992; Arnold, 1999). The list price choice is far from trivial as it must balance two opposite objectives: the maximization of the final sale price and the minimization of the time spent on the market (Merlo et al., 2015). This trade-off yields a complex problem for the seller, especially since it is likely that she has not enough information about demand (Anenberg, 2016). Indeed, houses are highly diffentiated assets and the seller is often unable to observe many previous transactions of similar properties nearby, making it harder for her to assess the demand she faces, and to decide of the optimal list price accordingly.

In this article, I investigate the home selling decision of imperfectly informed sellers. To do so, I build and estimate a single-agent dynamic search model of the housing market where the seller is uncertain and learns in a Bayesian way about demand. The estimated model yields insights on the optimal list pricing strategies and on how information frictions affect them. In particular I can estimate the value of additional information for the seller.

The core of the model mimics the main issue: each period, based on the information she has, the seller picks a price at which she lists her property in order to attract at least one potential buyer and eventually bargain and sell him her home. My main contribution is that I allow for several buyers on the market and make the seller learn about how many of them she expects to face. To do so, I split the demand side of the model into two components: the number of buyers present on the market and their valuations. Each buyer on the market decides to enter the bargaining game only if he expects to benefit from it, i.e. if his valuation is high enough compared to (a function of) the list price. Thus the buyer visits observed by the seller are endogenous and depend on the signal sent by the list price: ceteris paribus, a lower list price implies more visits.

This flexible specification of the demand helps disentangling low individual demand valuation from demand illiquidity. For example, even with an extremely low list price, the seller is not warranted to sell her home as there might simply not be any buyer on the market at this period (illiquid demand).

In addition, by introducing the possibility of several buyers, I can make sense of sales occurring above the list price and not only below or exactly equal to it. Since the list price serves as a commitment device for the seller, bargaining between her and one buyer can only result in a sale at a price inferior or exactly equal to the list price. Competition between several buyers may yield prices above the list price. I contribute to the bargaining literature (Rubinstein, 1982; Shaked and Sutton, 1984) by constructing and solving a sequential bargaining game between one seller and several buyers under complete information. The resulting bargaining rule determines the sale price in my model.

As data about sale prices are more accessible than data about the number of visits of any listed house, I model only one source of uncertainty for the seller: she knows the distribution of valuations and is only uncertain about the 'arrival of buyers on the market' modeled as a Poisson process. More precisely, the seller forms a prior belief about the arrival rate parameter. To update her belief, she never directly observes the latent number of buyers on the market, but only the number of buyers who decide to visit her house. This number of visits is a derivative process of the number of buyers present on the market and depends on the list price. It yields an original application of a Bayesian learning model where the decision variable (list price) influences the informational flow (by influencing the odds of observable entries) and thus the learning pace.

From a technical point of view, I contribute to the Bayesian learning literature by computing the updating rule for the prior belief about a Poisson process in the case where one only observes information coming from a known derivative of the original process he is willing to learn about.

Estimating the parameters of such a model where optimal strategies are time-dependent (non-stationary) requires detailed microdata. I use an original dataset of about 100,000 complete listing histories (dated initial list price and revisions up to final sale) of sold (with the help of a realtor) single-family homes in Pennsylvania between 2011 and 2014. I collected the data on 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 in the few existing other data (England 1995-1998: Merlo and Ortalo-Magne, 2004; Merlo et al., 2015; California 2007-2009: Anenberg, 2016). In particular, I observe that most sales occur at a price below the list price (overpricing in 74% of the cases), many at exactly the list price (15%) and the rest above it (underpricing). I also observe a well known puzzling fact about list prices: they are duration-dependent and generally adjust downwards throughout the listing process, even when market conditions seem stable (Salant, 1991).

In terms of results, my learning model (11 structural parameters only) matches the data closely. I estimate that an initially large positive expectation bias, progressively corrected via learning from the seller can explain the observed facts: in particular the decline of the list price and the distribution of time spent on market. It provides another explanation to the stylized facts with forward-looking rational agents (Merlo et al., 2015; Anenberg, 2016), in addition to other documented alternatives resorting to loss aversion and backward-looking agents (Genesove and Mayer, 2001) for example.

This paper also contributes to the literature on the role of overconfidence (Odean, 1998). 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. The value of additional information might even be negative as an overconfident seller can manage to obtain a better outcome (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. She also refuses offers more easily and sets higher list prices (implying less 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.

Related Literature: This paper builds upon the literature on dynamic search model of the housing market (Salant, 1991; Horowitz, 1992; Carrillo, 2012; Merlo et al., 2015; Anenberg, 2016, etc.). I try to explain several of the puzzling stylized facts about housing market which were first described in Merlo and Ortalo-Magne (2004). In particular I focus on the duration-dependent declining list price (first theorized by Salant, 1991) and on the relation between sale and list prices.

To address how optimal list price could decline over time, one needs to work in a non stationary finite horizon framework. Due to computational intensity and rich microdata requirements, only a few contemporaneous papers were able to estimate it: Merlo et al. (2015) and Anenberg (2016).

My work is close to the one of Anenberg (2016) who also builds a home selling model with imperfect information. He formulates and estimates a model where sellers are uncertain about the buyer valuations for their property and have 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 (declines, but also rare increases for example). The main difference with my model is that I allow for several potential buyers. Thanks to this, I can develop more extensively the bargaining side of the problem. For example it permits me to endogenously generate underpricing (list below sale price), while Anenberg (2016) is obliged to generate it via an exogenous probability parameter. This is particularly important as the phenomenom is not negligible in American data (11.42% of the time in my sample), compared to the English data of Merlo et al. (2015) for example (3.9%).

In another work, Merlo et al. (2015) successfully reproduce most data features without considering any information friction. To make sense of the optimal decline, they build and estimate a model with a rich time-dependent arrival probability function: exogenous changes in the arrival rate of buyers through time modify the seller optimal list price choice and make it can decline optimally. Similarly, they are able to fit the relation between sale and list prices using a rich parametric bidding function which depends on the number of matches/bidders. Overall they obtain a better fit than my model, at the expense of an heavily parametrized model, while I get a decent fit using only 9 structural parameters in my learning model.

The technical developments in this paper contribute to the bargaining literature by providing the solution of a sequential bargaining (Rubinstein, 1982) framework with one seller and N buyers. My proof is closely related to the one provided by Shaked and Sutton (1984): other buyers than the one with the highest valuation act almost as outside options for the seller.

This paper proceeds as follows. Section 2 describes the data. Section 3 develops a dynamic microsearch model of the home selling problem in which sellers learn. Section 4 details the estimation methods. 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 the American State of Pennsylvania, gathered on the American real estate platform *zillow.com*. The website is one of the leading online real estate market place. It is a real estate listing aggregator which gathers listings from real estate agents, from a partnership with MLS services and from private national companies (Century 21, Coldwell Banker and Sotheby's for example). As a consequence, I have one of the most exhaustive data about real estate transactions in Pennsylvania.

I focus on sales of single-family homes for which there is a complete history record of the transaction available. For each of these sales, my data include usual property attributes (square footage, lot size, number of beds, number of baths, year when the house was built, etc.). For about 33% of all transactions, a complete record of the last transaction history is available.As displayed in figure 1, this history contains all seller's decisions: the initial list price, its eventual adjustments through the sale process, potential intermediary listing removals and 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 5 months on market in 2013, after 4 list prices adjustments. It was finally sold at a price of $328000, slightly below the final list price of $329900.

In order to study list price dynamics and time on 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 allows home-owners to list their house on their own ('for sale by owner'). Unfortunately, these listings represent less than 1% of the observations. As a consequence, I cannot use my data to understand the seller's choice to resort or not to a real estate agent (as modeled by Salant, 1991), and I focus only on transactions in which a real estate agent intervened (as in Merlo et al., 2015 or in 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 which 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 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. 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 have already tried to sell his property (and failed) is different from a 'new' seller. In particular, these two types of sellers should be different in terms of the information they have about the market conditions (one of

[1]It 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 ones of properties which were relisted and sold afterwards.

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 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 eventual rejection of these offers by the sellers. Thus I am forced to keep the buyer's side of my model quite simple (contrary to Merlo et al., 2015 for example).

After selecting the observations and cleaning the data, I end up with 97 451 real estate transactions in Pennsylvania between August 2011 and July 2014. My data is described in the next section.

2.2 Summary statistics

Table 1: Summary Statistics

Variable	Mean	Std. Dev.	Min	Median	Max
Price and Timing:					
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				
Properties characteristics:					
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
Number of transactions	97451				
Number of census tracts	1806				

Table 1 and figure 2 present the summary statistics of my sample. The Pennsylvanian data exhibit common stylized facts to the ones observed in the literature and extensively detailed by Merlo and Ortalo-Magne (2004). First, prices are sticky and not often adjusted (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 close to 75% in another Zillow data set of properties on the United States' East Coast. I choose not to focus on the stickiness and not to model

Figure 2: Descriptive Statistics

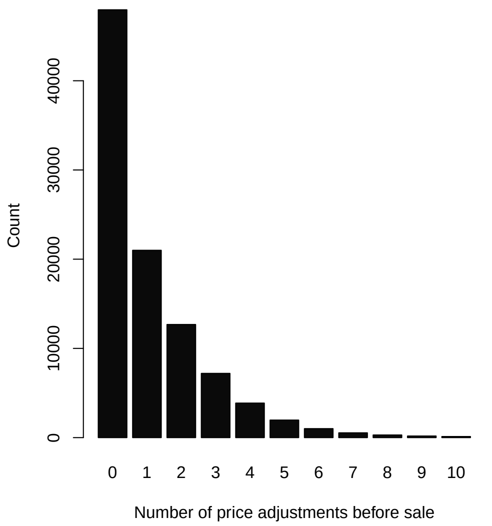

(a) Distribution of the number of price adjustments before sale

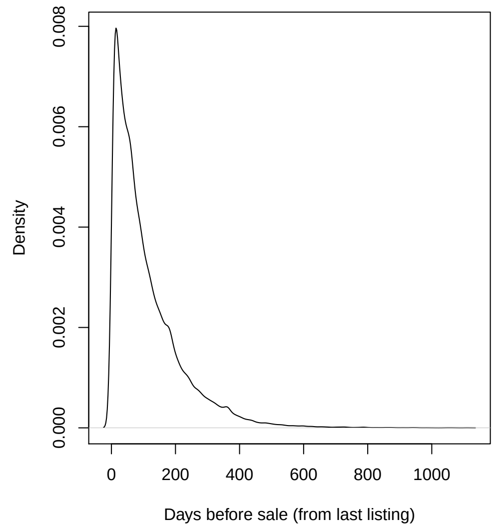

(b) Time on market distribution

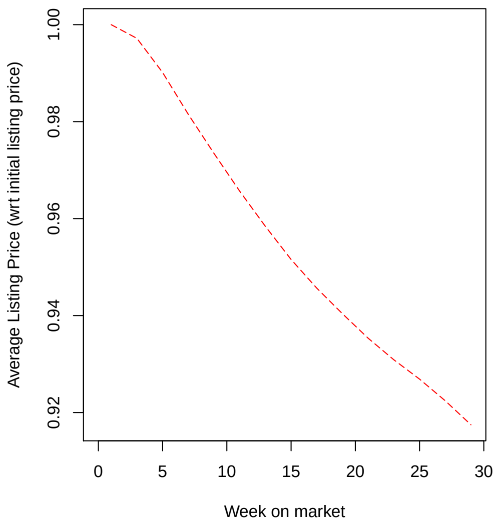

(c) list price dynamics (as ratio of initial list price)

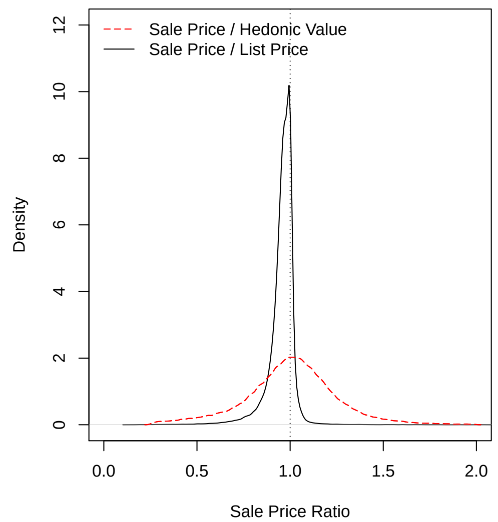

(d) Ratio Sale to list price distribution

it as including a menu cost would be too computationally costly and it has already been well explained by Merlo et al. (2015).

The average time on market is about 16 weeks (figure 2b). It is within the usual range: higher than the one observed by Merlo et al. (2015) in the UK (about 10 weeks), and slightly lower than the one 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 30^{th} 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 is considerably greater than in the English data (3.9% in Merlo et al., 2015): which is why my model emphasizes on endogenizing sales above the list price.

Unobserved Heterogeneity: 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 are not contained in the explanatory variables used in the hedonic estimation (number of bath, number of beds, living area, census tracts, etc.). I find that list price is a better predictor of the sale price than the hedonic fitted price, as it is well known in the literature (Horowitz, 1992; Merlo et al., 2015). To account for the extra information embedded in list price and unobserved by the econometrician, I do two things. First, my model includes some unobserved heterogeneity known by the seller. Second, to estimate the model I use mostly 'relative moments', normalized by the list price (for example sale price distribution relative to the final list price, or list price dynamics relative to initial list price) in order to soften the impact of this unobserved private information.

In the next section, I explain the model without the unobserved heterogeneity first, in order to simplify the exposition. I show how I include it at the end of the section.

3 Model

I model a discrete-time, finite horizon (26 periods of 2 weeks) problem of a homeowner deciding to sell her property with the objective of maximizing the sale price. As Anenberg (2016), I introduce uncertainty and Bayesian learning into this framework. To explain the previously described stylized facts, one of the main feature of the model is that the seller is imperfectly informed about the true demand. More precisely, she does not know the expected number of buyers in the market each period and 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 sale. The chosen optimal list price will balance a classic trade-off between high sale price and short time on 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 to accept an offer or not, knowing that if she refuses (or if she receives no offer), she will incur an 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. Then 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 $v^b_{(1)} < v^b_{(2)} < ... < v^b_{(n)}$). 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 1 seller with complete information have a unique subgame-perfect equilibrium outcome (cf Appendix A) which is my *bargaining rule*:

- if $v^s > v^b_{(n)}$ (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

$$p^S = max\Big\{v^b_{(n-1)}\ ,\ min\Big(p^L\ ,\ v^s + \frac{1}{2}(v^b_{(n)} - v^s)\Big)\Big\}$$

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

In words, trade only occurs if there is gains from trade ($v^b_{(n)} \geq v^s$). 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 list price, bilateral sequential bargaining between the seller and buyer$_{(n)}$ would yield the classic Rubinstein (1982) outcome that they share the 'transaction gains'. The seller gets a portion ϕ of $(v^b_{(n)} - v^s)$ while the buyer gets the rest (the portion $1 - \phi$). Thus the sale price would be $p^S = v^s + \phi(v^b_{(n)} - v^s)$. 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 $p^S = min\Big(p^L\ ,\ v^s + \frac{1}{2}(v^b_{(n)} - v^s)\Big)$.

Now, if another buyer (the second highest valuation buyer, denoted buyer$_{(n-1)}$) is present on the market,

[2] To end up with an equal sharing of the pie 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 (cf 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 implicitely in Anenberg (2016).

his valuation $v^b_{(n-1)}$ can serve as an outside option for the seller.[3] If $v^b_{(n-1)}$ 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 $p^S = v^b_{(n-1)}$ 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 ($v^s + 1/2(v^b_{(n)} - v^s)$), at the list price ($p^L$), and above it (in case of competition, with $v^b_{(n-1)}$ which can be higher than p^L sometimes). The challenge being 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 to compute what would be the hypothetical bargaining outcomes (sale or not, and eventual sale price) for any scenario. As a consequence, for a given number of buyers, if she knows the buyers' valuations distribution, the seller is able to build expectation about the bargaining outcomes (using order statistics of the highest and second highest valuation for the given number of buyer). If she also has an idea of the distribution of 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 splitted in two main components: the number of buyers on the market during the period ($N^{market}_t \sim Poisson(\lambda)$) 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 by the property (N^{market}_t) is drawn from a $Poisson(\lambda)$ distribution.The rate of arrival λ is a key demand parameter in the model. The seller does not know about λ, she forms an initial belief about it that she will update using Bayesian learning.

[3] Notice that only the second highest valuation can impact the sale price, ceteris paribus the other buyers' valuations are irrelevant.

The reservation values of every buyers 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 implicitely assume 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 $200,000$\$. This way I build a single representative problem for every sellers, independently of the 'quality' of their properties. Then, under the linear homogeneity assumption, I can compare data counterparts to the model price outcomes: the data prices normalized by their hedonic values.

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

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 bargain in the model. Before inspecting, the buyer observes a detailed ad about the property on the listing website and he 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 it is beneficial

[4] An 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'. And because of its cost (going to the bank and dealing with all financial details), buyers need to decide whether to 'inspect' (make an offer) or not. If one choose to enter the bargaining, he learns about the seller and his competitors values. The final outcome is determined via the bargaining rule.

for him to enter or not, 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 $g(x) = a_0 + a_1 x$ with $0 < a_1 < 1$. It yields the following simple *inspection/entry rule* that all buyers follow:

$$\text{a buyer inspects the property if } v^b > g(p^L) = a_0 + a_1 p^L$$

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 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 $Pr(v^b > g(p^L))$ 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 higher taste for her property, leading to a higher expected sale price, but it 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 market. As staying on the market is costly for the seller (she has to keep his home tidy, spend time for potential visits, etc.), the optimal list price, which maximizes the seller expected profit from sale, balances two opposite objectives: short time on market and high sale price.

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

3.3 Seller's information and learning

When she decides to set her list price, the seller 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 $Poisson(\lambda)$ distribution. However, she has imperfect information about the demand since she *does not know the value of* λ.

In period 0 (at the start of the listing), the seller forms an initial belief about it: $\lambda \sim Gamma(\alpha_0, \beta_0)$,

[5] Buyers are really 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 one period. They do not have time to gather information, thus they build really 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 buyer believed functional form $g()$ itself. I show in the results 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).

that she will update each period via Bayesian learning rules.

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, $N_t \sim Poisson\big(\lambda \; Pr(v^b > g(p_t^L))\big)$: to learn about λ, the seller does not directly observe the latent process determined solely by λ, but another process which is more or less close to it depending on the choice of list price. Nonetheless, computing the posterior distribution remains quite simple in this case. For simplicity, let's 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*:

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

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, remark that when the list price is so high that $Pr(v^b > g(p^L)) \to 0$, 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 $Pr(v^b > g(p^L)) = 0$, 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, $Pr(v^b > g(p^L))$ 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 $Pr(v^b > g(p^L))$). Up to the opposite extreme scenario where p^L is so small that $Pr(v^b > g(p^L)) \to 1$, 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

[6] In this model λ is fixed and does not vary with time. But 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).

inspection process' closer to the latent 'number on market process' and allows to learn faster about the demand parameter determining this number of buyers on the market.

3.4 Seller's optimization problem

Figure 3: Timeline of events in the model

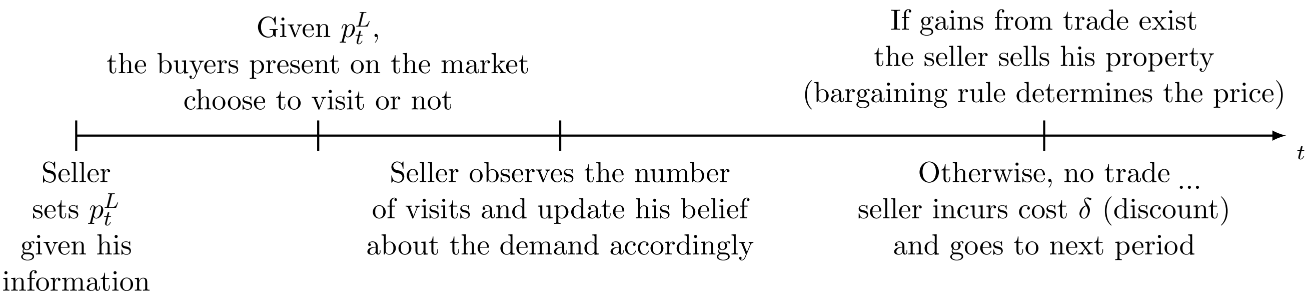

The timing of the model is summarized in figure 3. At the start of each period, the seller sets an optimal list price p_t^L in order to maximize her expected profit from sale given her information set. Her information set consists of her belief about λ at the start of the period (which can be summarized 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, 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 update her belief about λ according to the Bayesian learning rule: it pins down 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 next period where she repeats the same process, starting with her updated belief.

I repeat this game over a finite horizon of 52 weeks (26 periods): if at the end of this year, the seller did not sell her property, the simulation ends without sale.

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

[7] The 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 which is build based on her expected gain (in the present or future): thus it only impacts her choice to leave the table in the bargaining.

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:

$$v^s(\Omega_t) = \underset{p_t^L}{max} \sum_{k=0}^{\infty} \mathbb{E}\Big[Pr(N_t = k \; |p_t^L)\Big|\Omega_t\Big] \; \mathbb{E}\Big[\Pi^s(N_t = k, \; p_t^L, \; \Omega_t) \; \Big]$$

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

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

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

and the corresponding profit function depends on the number of inspections and known updating in the case where this number of inspection indeed happens (i.e. $v^s(\Omega_{t+1})$ instead of Ω_t), as defined below:[8]

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

with $\Omega_{t+1} = \Big(\alpha_t + N_t, \; \beta_t + Pr(v^b > g(p_t^L))\Big)$. 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 matter potentially 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 $i < j \in 1, 2, ..., n$, $\forall x < y \in \mathbb{R}$ 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.

3.5 Seller's heterogeneity

I normalize every prices by their hedonic values in order to be able to compare the data with a single representative seller problem in the model. However, as explained when I described the data, the list price embeds some extra private information about the sale price that is not captured by the hedonic

[8] For any number of visits that occur when the seller chose 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 she would bargain with in each specific entry case $(v^s(\Omega_{t+1}))$. While the probability of observing each specific number of entries are computed using the period starting belief Ω_t.

values. Thus the data exhibit unobserved heterogeneity for the econometrician. In order to account for this, I also add unobserved heterogeneity within the model, by considering two types of sellers:

$$\begin{cases} \text{Seller facing 'normal' demand:} & v^b \sim \mathcal{LN}(\mu, \sigma) & \text{with proportion } 1-\pi \\ & \text{Initial belief: } \lambda \sim Gamma(\alpha_0, \beta_0)) & \\ \text{Seller facing high demand:} & v_h^b := \kappa v^b \sim \mathcal{LN}\Big(\underbrace{\mu + log(\kappa)}_{=\mu'}, \sigma\Big) & \text{with proportion } \pi \\ & \text{Initial belief: } \lambda \sim Gamma(\alpha_0^h, \beta_0)) & \end{cases}$$

In words, there are some sellers (proportion π) who have private information that their properties is of higher quality (for some reason unobserved by the econometrician). I refer to them as 'high demand sellers'. These sellers know they are facing a higher demand than the others: κv^b instead of v^b, where κ is just a scaling parameter > 1. Essentially, the problem is exactly the same as the one described in the previous section, the presence of κ simply shifts the mean of the valuation distribution upwards to $\mu' = \mu + log(\kappa) > \mu$. Otherwise they face exactly the same problem. In particular they face the same unknown λ. However, as they represent a different population, I allow them to have a different initial belief about it, via the α parameter: $\lambda \sim Gamma(\alpha_0^h, \beta_0)$ instead of $Gamma(\alpha_0, \beta_0)$.
The other sellers (proportion $1 - \pi$) face exactly the problem described before (with μ and α_0).

When he observes the model results, the econometrician has no variable to know if a seller faced a high demand, as it is the case in the data. This way list prices embed extra private information about the sale prices and are closely correlated to it for unobserved reasons in the model too.
The extra heterogeneity within the model allows for a better match of the data with the estimation.[9]

4 Estimation method and identification

The structural parameters that I want to estimate are: $(\lambda, \mu, \sigma, a_0, a_1, \delta, \kappa, \pi, \alpha_0, \alpha_0^h, \beta_0)$. Denote this 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' behaviour will be the closest to the observed sellers' behaviour. To do so, I select features from the empirical data

[9]For computational burden reasons, I can only include two types of demands, and not a wider specter of them. First, including more than two types requires to add at least 2 structural parameters (the scale, the proportion in the population and eventually the initial belief difference) for each of the new type which is too costly in terms of estimation. One could solve it by defining a known distribution for κ (implying a small additional parameter burden). However, the main limitation comes from the value function computational burden. Indeed computing a value function $v^s(\Omega_t, \mu, \sigma, a_0, a_1, \delta)$ is long, and I would need to compute one for every scale (i.e. every values of μ). This is already too long with only a single κ.

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 eleven unknown parameters by minimizing a distance function between empirical and simulated moments such that:

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

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 11 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 market, the distribution of sale price and the list price dynamics. Most of the price moments are relative to the initial list price. It allows me to reduce issues caused by the 'scaling' of the problem or to soften the impact unobserved heterogeneity not accounted for in hedonic value estimation. I only have 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 should allow to identify the parameters. Intuition about identification is non trivial as most parameters influence several features of the model simultaneously. For example all parameters (except λ) will directly influence the list price choice, and even λ will influence it indirectly through the updating of the initial belief. A rough intuition can be given as follows: μ and σ are identified by the distribution of sale price, especially the ones above the list price (since in this case, following the model, one observes $p^s = v^b_{(n-1)}$ directly). Along them λ pins down the time on market. a_0 and a_1 are identified by the sale price distribution, in particular its lower tail. δ adjusts to pin down the remaining part of the sale price distribution correctly, as well as the list price value. By adding some heterogeneity, κ and π captures some of the unobserved heterogeneity present in the data. It helps to match the normalized mean sale price and initial list price while also improving the match of other moments. Finally, the three initial seller beliefs parameters pins down the initial list price (mixed with π for the averaging), and alongside with the true value of λ they determine the dynamic adjustment.

[10]Recall 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 sellers (independent of the 'quality' of the property). Hence I can run iid simulations of the representative problem and compare it to normalized real data.

5 Results

5.1 Value function and optimal list price

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

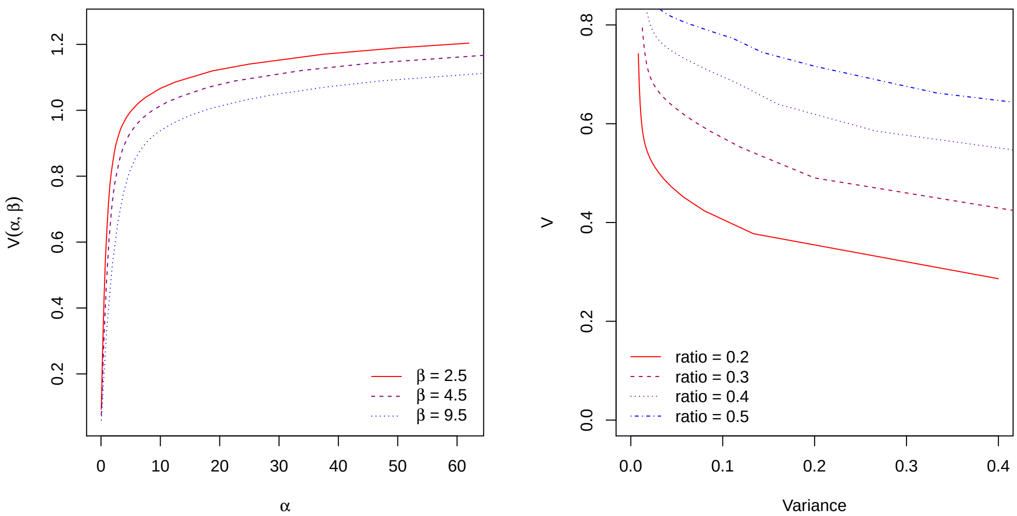

(a) Value as function of α for fixed β

(b) Value as function of variance α/β^2 for fixed expectation α/β.

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

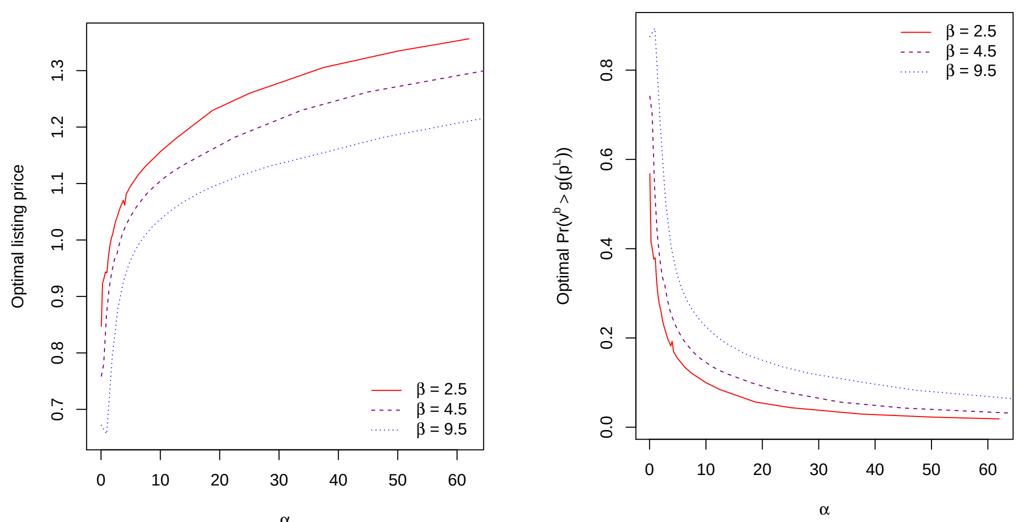

(a) p^L as function of α for fixed β

(b) $Pr\big(v^b > g(p^L)\big)$ as function of α for fixed β

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

I obtain intuitive results for the value function (figure 4) and optimal list price (figure 5). First, the value increases with the expected number of buyers on the market. Second, at fixed average belief, the less uncertain the seller is (i.e. the smaller variance of the belief), the highest value she obtains. 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 $\lambda Pr\big(v^b > g(p^L)\big)$.

5.2 Parameters and Moments

Table 2: Parameter estimates of Structural Model

Parameter	Description	Estimate	Std. Errors
δ	Subjective discount factor (and cost of listing)	0.9819	0.00021
Demand parameters	$N^{market} \sim Pois(\lambda), \quad v^b \sim \mathcal{LN}(\mu, \sigma)$ Proportion π of high demand sellers facing κv^b		
μ	Mean of buyer valuation	-0.1126	0.00058
σ	Standard deviation of buyer valuation	0.1433	0.00035
λ	Expected number of buyers on the market (2 weeks)	0.3936	0.00125
κ	Scaling of valuations in case of high demand	1.2483	0.00128
π	Proportion of people facing high demand	0.2635	0.00233
Belief	$\lambda \sim Gamma(\alpha_0, \beta_0)$ or $\lambda \sim Gamma(\alpha_0^h, \beta_0)$		
α_0	α prior belief distribution	3.95	0.02883
α_0^h	α prior belief distribution (high demand)	7.70	0.06547
β_0	β prior belief distribution	5.50	0.02091
Inspection rule	buyer inspects if $v^b > a_0 + a_1 p^L$		
a_0	Buyer's conjecture about seller reservation value: constant	0.3191	0.00078
a_1	Buyer's conjecture about seller reservation value: slope	0.6525	0.00071

Table 2 reports the parameter estimates. The sellers are clearly overestimating the demand (λ). The 'normal' sellers start with an initial belief $\lambda \sim Gamma(\alpha_0 = 3.95, \beta_0 = 5.5)$ and thus expects lambda to be almost twice as high as what it truly is ($\alpha_0/\beta_0 = 0.718 > \lambda = 0.3936$). The high demand sellers are more confident about the demand than the normal sellers. Their expectation bias is larger ($\alpha_0^h/\beta_0 = 1.40 > \lambda = 0.3936$) about 3.5 times the true λ. My results suggests that the progressive correction of the initial (large) positive expectation bias via learning, added to a natural selection/survivor effect (the one who stays on the market are likely to observe less entries and thus to have decreasing belief) are the reason for decreasing list price.

Another thing worth noting is the screening of buyers implicitely done by the seller at the optimal parameters. Take the seller facing normal demand $v^b \sim \mathcal{LN}(\mu = -0.1126, \ \sigma = 0.1433)$ for example. These sellers set an initial listing price of 0.9964 (not displayed in the results table). In this case, $g(p^L) = 0.9693$, $Pr(v^b > g(p^L)) = 28.49\%$ and the true rate of inspection $\lambda Pr(v^b > g(p^L))$ is equal to 11.21%. The seller aims for high quantiles of the demand (top 28.49%), but she actually does so 'by mistake' because she expects a higher rate of inspections than what it truly is (since she overestimates λ). As she learns about λ (decreasing expectation), the seller corrects her 'mistake' progressively by decreasing her list price, and thus by increasing $Pr(v^b > g(p^L))$ in order to balance the inspection rate. The pattern is exactly the same for the high demand sellers.

Figure 6: Actual and Simulated moments

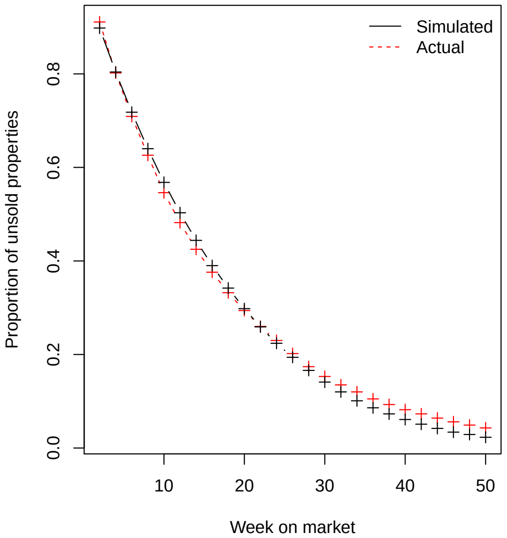

(a) Proportion of unsold properties as a function of time on market

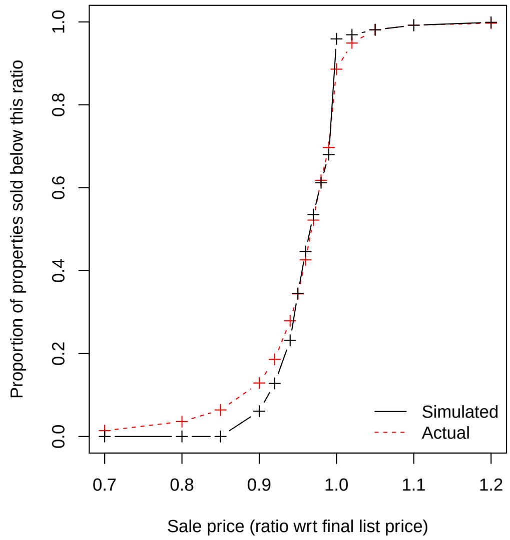

(b) Distribution of sale prices (ratio to final list price)

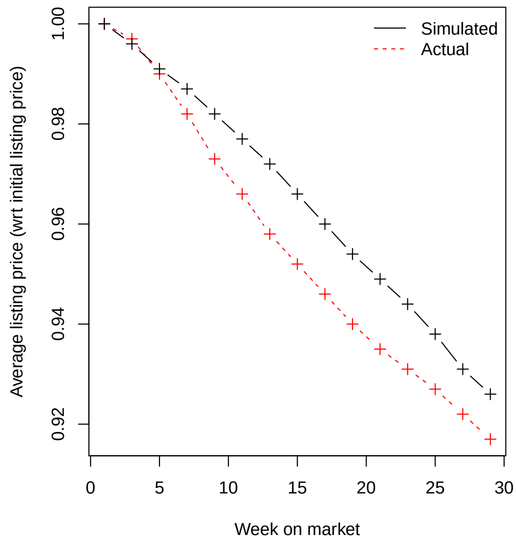

(c) Average list price as a function of time on market (ratio to initial list price)

Table 3: Actual and Simulated moments

Moment	Actual	Simulated
Mean sale price	1.008	1.025
Mean ratio sale/final list price	0.955	0.966
Mean initial list price	1.107	1.106
% of accepted offers equal to list price	0.15	0.222
% of accepted offers below list price	0.734	0.737
Mean week on the market (knowing that <52 weeks)	14.817	15.71

Table 3 and figure 6 illustrate how the model matches the moments. Overall, for a small number of parameters (11), it fits the data correctly. The list price dynamics and the distribution of the time on market are well fit. 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 4.1% in the simulation against 11.6% in the data). This is because, even with the splitted demand, the model is still unable to match the time on market distribution and the sales above the list price at the same time.

6 Value of information

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 λ. I denote $v^s_{PI}(\lambda)$ the value of this perfectly informed seller.

I compare this benchmark to the expected value of an imperfectly informed (denoted II) seller with belief (α, β) about λ, i.e. $v^s\big(\lambda \sim Gamma(\alpha, \beta)\big)$ (which corresponds to the value function, simply denoted $v^s(\alpha, \beta)$, or $v^s(\Omega)$ earlier). The comparisons in this section are based on the model for a seller facing 'normal' demand computed at the corresponding estimated structural parameters. I only vary the initial beliefs parameters (by varying α for a fixed $\beta = 5.5$ $(= \beta_0)$, or by varying both parameters for a fixed ratio $\alpha/\beta = \lambda$ in figure 7).

Figure 7: Convergence of imperfect information to perfect information value

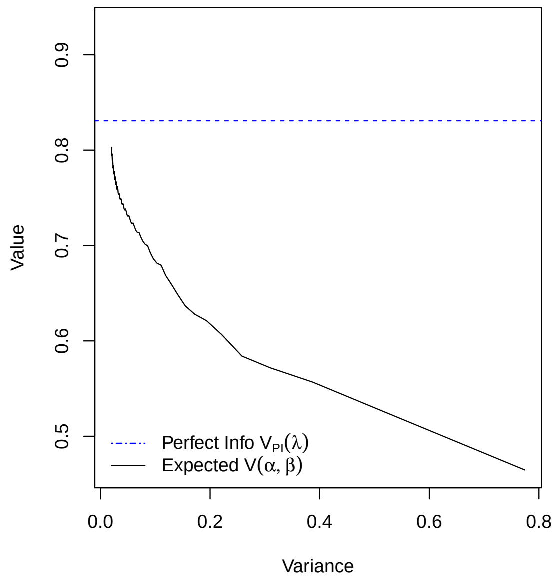

Comparison of value as a function of changes in variance, for fixed λ (=0.3936) and fixed mean belief (ratio α/β) equal to λ.

First, notice from figure 7 that when the belief of the II agent converges to λ being equal to its true value, her value also converges to the PI value. This makes sense since the II seller is optimizing a problem closer and closer to the PI seller in this case.

Figure 8: Value of information: sale price and TOM

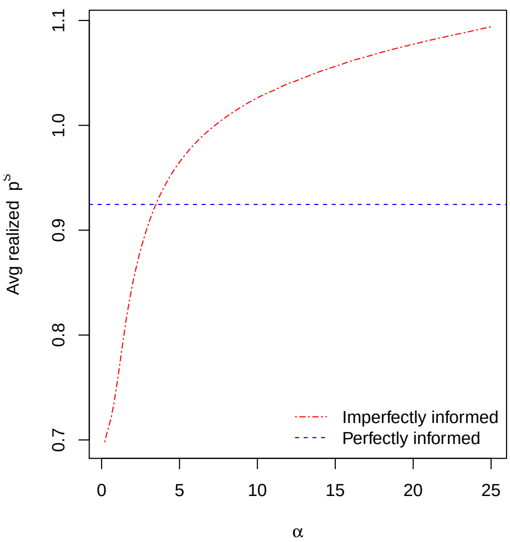

(a) Average obtained p^S (not discounted) as a function of initial belief (varying α with fixed $\beta = 5.5$), for fixed λ (=0.3936)

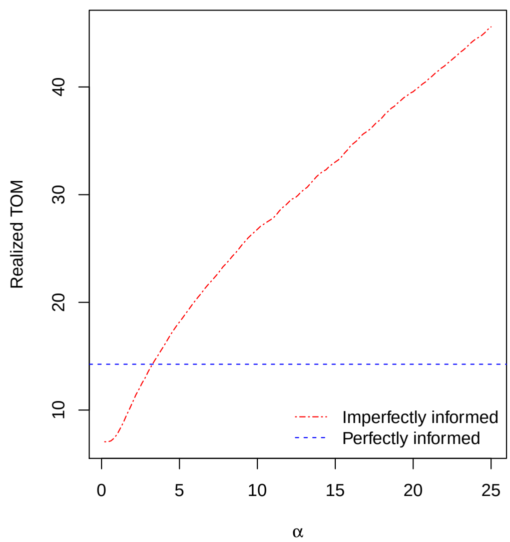

(b) Average time spent on market (TOM) as a function of initial belief (varying α with fixed $\beta = 5.5$), for fixed λ (=0.3936)

Figure 9: Value of information: information paradox

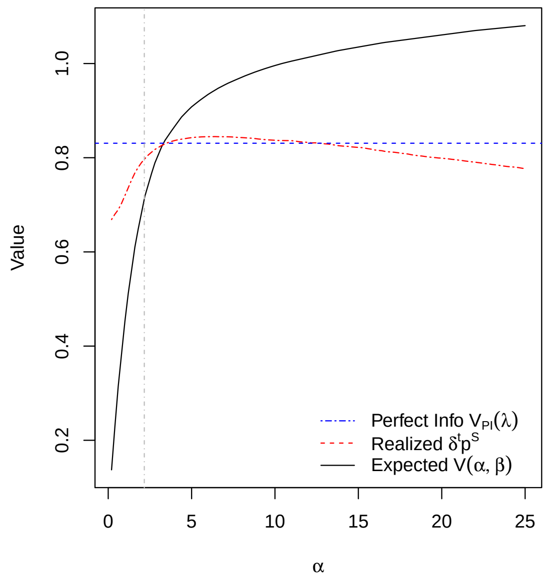

Value as a function of initial belief (varying α with fixed $\beta = 5.5$), for fixed λ (=0.3936).

More interestingly, I also compare the PI benchmark to the true realized outcomes of an agent who starts with a given initial belief (α, β): i.e. for a given initial belief and a true λ, I simulate the complete selling process (with corresponding dynamic belief updating) and see what price p^S the seller obtains on average (discounted from the initial period, i.e. $\delta^{\text{time on market}} p^S$). This average discounted price is the simulated counterpart to the theoretical seller value (v^s) thus I can compare them directly.
Notice that this obtained realized average discounted price is truly different from the seller's expectation: as already mentioned, $v^s(\Omega)$ only represents the expectation of the seller based on her belief, and does not depend on λ at all. If the belief is completely wrong, the seller will simply overestimate her initial expected value from the selling process, and shall update it progressively as she will learn from observing the market reactions. As a consequence, to get an idea of the value of information, I need to compare the perfectly informed benchmark valuations to counterparts realized outcomes of imperfectly informed sellers (instead of to pure expectations based on the belief only). It will allow to see how far imperfectly informed sellers end up from the perfect information benchmark depending on their initial misinformation level.

From figure 8, I observe that the more overoptimistic the starting initial belief (higher α for fixed β), the higher the obtained average sale price (figure 8a), but also the longer the average time spent on market (figure 8b). Notice that the average sale price obtained with a wrong overoptimistic belief might be higher than the average sale price obtained by a perfectly informed agent. However, this might not necessarily yield a higher value (discounted price) for the II agent as she also spends more time on market, which means that her outcome is discounted more strongly.

Figure 9 shows how it translates in terms of value... And it yields an *informational paradox*. Indeed, there is an area where the overoptimistic misinformed agent can perform better than the perfectly informed one in this model, or in other words, 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 with respect to the true demand parameter λ. This means that having a better information about λ is not necessarily beneficial for the seller, the value of additional information might be negative!

This seems quite puzzling 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* $v^s(\alpha, \beta)$. It gives her a better 'bargaining position' (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, $v^s + 0.5(v^b_{(n)} - v^s) = 0.5v^s + 0.5v^b_{(n)}$ is also higher. There is some level of overconfidence where the overconfident seller obtains sale price sufficiently high to offset the longer time she spends on the market, yielding a better outcome overall. At some point, the gains in bargaining position are offset by a too large misoptimization (too long time spent on market), and being too overoptimistic yields a lower value than being perfectly informed.

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 infinitely more patient than the buyers in the bargaining game. This way she has all bargaining power and can extract $v^b_{(n)}$ without splitting the pie, such that v^s disappear from the bargaining rule. And in this case, the information paradox disappears: the perfectly informed agent always performs better.

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.

My estimated model suggests that home-sellers exhibit initial over-optimistic expectation about the demand that they progressively correct by learning throughout the selling process. This initial expectation bias is a key feature of the model and I succeed in generating a realistic simulated list price dynamics and distribution of time spent on market thanks to it.

My work also highlights a paradox that the value of information is not necessarily positive. Indeed, by being overconfident about the demand, the seller overestimates her reservation value and has a stronger bargaining position. This allows her to extract a higher sale price, which can compensate her misoptimation and longer time spent on the market (with respect to the perfect information benchmark).

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

Appendix A Bargaining rule

I study a simplified version of the problem with one seller and only 2 buyers and with normalized valuations (as described right after). This generalizes easily to N buyers with ordered valuations: only the two buyers with the two highest valuations matter with complete information.
The proof is inspired from Shaked and Sutton (1984) and Binmore et al. (1989), i.e. bilateral bargaining with outside option (for the player playing second). The problem setup and thus the result is also close to the 'auctioning model' in Binmore et al. (1992).[11]

A.1 The Problem: 1 Seller - 2 Buyers

There is one seller of an indivisible good, with reservation value $v_s = 0$. The seller can sell it to one of two buyers H and L (high and low) with valuations $v_H = 1 > v_L = v$.[12] The three agents have a common discount factor ρ. The period length is τ. Denote for simplicity $\delta = \rho^\tau$. If the good is sold at price p after t periods, the seller's payoff is $p\delta^t$, the successful buyer's payoff is $(v_b - p)\delta^t$ and the losing buyer gets zero. Information is complete.
The *timeline* of the problem is as follow:

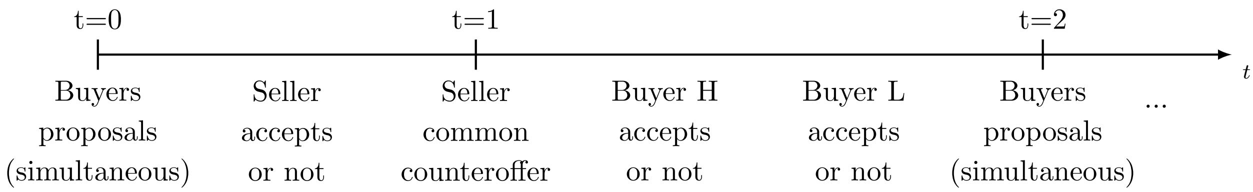

In period $t = 0$, each buyer simultaneously makes a proposal to the seller. She may accept one of these offer or rejet both. If the seller accept one of the offer, the two players trade and the game end. For simplicity, if the seller wants to accept an offer that both buyers made (i.e. in case of tie), the tie breaking rule is that the seller will opt for buyer H.
If the seller reject both offers, there is a delay τ and they go to the next bargaining period. The seller then make a common counteroffer to both buyers. The highest valuation buyer either accept or reject it. If he rejects it, the low valuation buyer can then choose to accept it or not. If both buyers 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.

[11]Except for a different timeline (buyers start for me, which change the result), and some notation/normalization changes. However no proof is provided in the book, and the references given to find the proofs are unavailable/unpublished or do not contained the proof at all.
[12]I normalize $v_H = 1$ so that the 'size of the cake' to share is one as usual. I denote v_L as v for simplicity.

A.2 Solution

This game always have 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 force buyer H to pay a higher price than if buyer L was absent.

A.3 Proof

As in 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 e.g.). Let's 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.[13]

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

$$m_b \geq v - max\{\delta M_s,\ s\} \tag{1}$$

$$v - M_b \geq max\{\delta m_s,\ s\} \tag{2}$$

$$m_s \geq v - \delta M_b \tag{3}$$

$$v - M_s \geq \delta m_b \tag{4}$$

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 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 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

[13]Obviously 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.

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:

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

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

  $1 - m_s \leq \delta M_b \leq \delta - \delta^2 m_s$, thus $1 - m_s \leq \delta - \delta^2 m_s$, 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 $\delta m_s = \delta/(1+\delta) \geq s$ $\iff s \leq \delta/(1+\delta)$. [14]

- 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.

  As a consequence 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:

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

[14] Intuitively, 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.
Let's 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 $\delta - \delta^2 M_b$). And 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 $\delta M_s = 1-\delta(1-s) \leq s \iff s \leq \delta/(1+\delta)$.

So, if subgame-perfect equilibria exist, they generate a unique subgame-perfect equilibrium outcome. And 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 allow 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.[15]

[15] Even 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. And 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 we could have expected, in equilibrium the low buyer will bid his valuation $s = v$.Which yield 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 just consider that in the proof here, 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, afterwards, 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).

The problem is as follows:

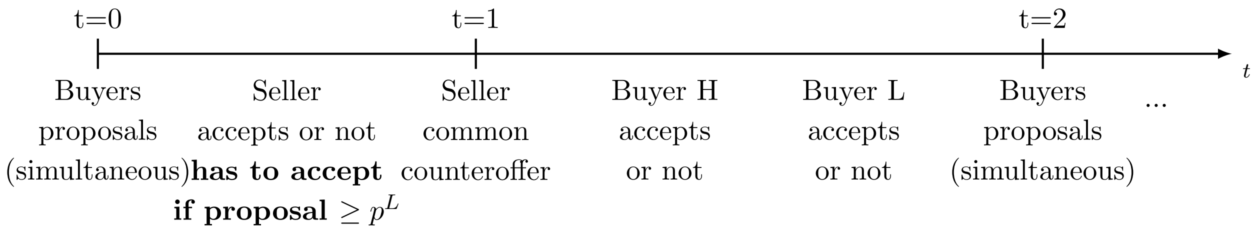

As a consequence, the list price only affects the buyer choice (in the first period), not the seller's (who only endure it). From the point of view of the seller, either she receives an offer higher than p^L and then she cannot 'play', or, she can play the game as before.
From the point of view of the buyer H, his choice is just to choose between offering p^L (or more) or bargaining with the seller as usual (i.e. entering the classic game). Thus it will be quite simple: if he gets more by offering just p^L, he will do that, otherwise he won't. Buyer H basically has a choice to resort to an 'outside option' before the game even starts. The only trick is that buyer L is still present and can also offer more than p^L: thus we still have a competition between H and L, even if H wants to bid more than p^L. It means that we do not have a classic outside option $= p^L$ for buyer H, but instead an outside option $= max\{p^L, v\}$: because recall that there is no equilibrium where H has to pay less than v, thus if $v \geq p^L$, then he will pay v instead of just p^L immediately (and obviously we have no concern about any bargaining since $v \geq p^L$ implies that the seller must accept immediately).
Basically we only have bargaining when $p^L > max\{\delta/(1+\delta), v\}$ now.

The equilibrium outcome is still that the seller accepts immediately (but sometimes because she has to accept now) the initial offer from the buyer H at price p, where p is now defined as follows:

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

$$= max\left\{ v,\ min\Big(\delta/(1+\delta),\ p^L\Big) \right\}$$

Now, if we assume that $\delta \to 1$ (either because the bargaining period $\tau \to 0$ or because both individual discount patience parameter $\rho \to 1$), we get an equal share of the pie $\delta/(1+\delta) = 0.5$. Moreover, if we 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	1.025
Mean ratio sale/final list price	0.955	0.966
Mean initial list price	1.107	1.106
% of accepted offers equal to list price	0.15	0.222
% of accepted offers below list price	0.734	0.737
Mean week on the market (knowing that <52 weeks)	14.817	15.71
% unsold after 2 weeks	0.911	0.898
% unsold after 4 weeks	0.802	0.804
% unsold after 6 weeks	0.709	0.718
% unsold after 8 weeks	0.626	0.64
% unsold after 10 weeks	0.546	0.568
% unsold after 12 weeks	0.482	0.503
% unsold after 14 weeks	0.425	0.444
% unsold after 16 weeks	0.376	0.39
% unsold after 18 weeks	0.332	0.342
% unsold after 20 weeks	0.294	0.298
% unsold after 22 weeks	0.26	0.259
% unsold after 24 weeks	0.23	0.224
% unsold after 26 weeks	0.202	0.194
% unsold after 28 weeks	0.174	0.166
% unsold after 30 weeks	0.153	0.141
% unsold after 32 weeks	0.135	0.12
% unsold after 34 weeks	0.12	0.101
% unsold after 36 weeks	0.105	0.086
% unsold after 38 weeks	0.093	0.073
% unsold after 40 weeks	0.082	0.061
% unsold after 42 weeks	0.073	0.051
% unsold after 44 weeks	0.064	0.042
% unsold after 46 weeks	0.056	0.034
% unsold after 48 weeks	0.049	0.029
% unsold after 50 weeks	0.043	0.023
Mean list price in week 3	0.997	0.996
Mean list price in week 5	0.99	0.991
Mean list price in week 7	0.982	0.987
Mean list price in week 9	0.973	0.982
Mean list price in week 11	0.966	0.977
Mean list price in week 13	0.958	0.972
Mean list price in week 15	0.952	0.966
Mean list price in week 17	0.946	0.96
Mean list price in week 19	0.94	0.954

Mean list price in week 21	0.935	0.949
Mean list price in week 23	0.931	0.944
Mean list price in week 25	0.927	0.938
Mean list price in week 27	0.922	0.931
Mean list price in week 29	0.917	0.926
% sales/list <0.70	0.014	0
% sales/list <0.80	0.036	0
% sales/list <0.85	0.064	0
% sales/list <0.90	0.129	0.061
% sales/list <0.92	0.186	0.128
% sales/list <0.94	0.279	0.232
% sales/list <0.95	0.344	0.345
% sales/list <0.96	0.426	0.446
% sales/list <0.97	0.522	0.535
% sales/list <0.98	0.618	0.612
% sales/list <0.99	0.697	0.68
% sales/list <1.00	0.886	0.959
% sales/list <1.02	0.949	0.969
% sales/list <1.05	0.981	0.981
% sales/list <1.10	0.992	0.992
% sales/list <1.20	0.997	0.999

References

Anenberg, Elliot (2016). “Information frictions and housing market dynamics”. In: *International Economic Review* 57.2.

Arnold, Michael (1999). “Search, Bargaining and Optimal Asking Prices”. In: *Real Estate Economics* 27.3, pp. 453–481. URL: http://EconPapers.repec.org/RePEc:bla:reesec:v:27:y:1999:i:3:p:453-481.

Binmore, Ken, Avner Shaked, and John Sutton (1989). “An Outside Option Experiment”. In: *The Quarterly Journal of Economics* 104.4, pp. 753–770. ISSN: 00335533, 15314650. URL: http://www.jstor.org/stable/2937866.

Binmore, Ken, Martin J. Osborne, and Ariel Rubinstein (1992). “Chapter 7 Noncooperative models of bargaining”. In: vol. 1. Handbook of Game Theory with Economic Applications. Elsevier, pp. 179 –225. DOI: http://dx.doi.org/10.1016/S1574-0005(05)80010-4. URL: //www.sciencedirect.com/science/article/pii/S1574000505800104.

Carrillo, Paul E. (2012). “An Empirical Stationary Equilibrium Search Model Of The Housing Market”. In: *International Economic Review* 53.1, pp. 203–234.

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

Horowitz, Joel (1992). “The Role of the List Price in Housing Markets: Theory and an Econometric Model”. In: *Journal of Applied Econometrics* 7.2, pp. 115–29.

Merlo, Antonio and Francois Ortalo-Magne (2004). “Bargaining over residential real estate: evidence from England”. In: *Journal of Urban Economics* 56.2, pp. 192–216.

Merlo, Antonio, Franois Ortalo-Magne, and John Rust (2015). “The Home Selling Problem: Theory And Evidence”. In: *International Economic Review* 56, pp. 457–484.

Odean, Terrance (1998). “Volume, Volatility, Price, and Profit When All Traders Are Above Average”. In: *The Journal of Finance* 53.6, pp. 1887–1934. ISSN: 1540-6261. DOI: 10.1111/0022-1082.00078. URL: http://dx.doi.org/10.1111/0022-1082.00078.

Price, W. L. (1983). “Global optimization by controlled random search”. In: *Journal of Optimization Theory and Applications* 40.3, pp. 333–348. ISSN: 1573-2878. DOI: 10.1007/BF00933504. URL: https://doi.org/10.1007/BF00933504.

Rubinstein, Ariel (1982). “Perfect Equilibrium in a Bargaining Model”. In: *Econometrica* 50.1, pp. 97–109. ISSN: 00129682, 14680262. URL: http://www.jstor.org/stable/1912531.

Salant, Stephen (1991). “For Sale by Owner: When to Use a Broker and How to Price the House”. In: *The Journal of Real Estate Finance and Economics* 4.2, pp. 157–73.

Shaked, Avner and John Sutton (1984). “Involuntary Unemployment as a Perfect Equilibrium in a Bargaining Model”. In: *Econometrica* 52.6, pp. 1351–64. URL: http://EconPapers.repec.org/RePEc:ecm:emetrp:v:52:y:1984:i:6:p:1351-64.