Public Information and Sequential Learning in a Bilateral Market*

Nikita Kotsenko†

19.8.2021

I study the effects of improved public information on equilibrium welfare and price dispersion, providing sufficient conditions for negative and positive effects. Public information affects welfare by reducing excessive (though rational) pessimism induced by sequential learning. Reduced pessimism results in fewer agents withdrawing from the market prematurely (compared with the full information benchmark), which increases own ex-post utility in expectation but may also cause a congestion externality. I show that either effect can dominate. Observed search duration on the short side of the market is an indicator of the welfare effects of public information. The context is a search, matching and bargaining market with uncertainty about the meeting probability on both sides. Fully rational, ex-ante identical participants gather private information endogenously through costly search. Full trade is the unique perfect Bayesian equilibrium under general conditions. Full trade implies learning terminates following a positive but not following a negative signal, which results in a declining belief path and in declining reservation prices during search. The results hold for any prior distribution where a more precise public signal slows learning.

*This work has benefited from the advice of too many to mention. In particular, my advisor David Genesove, my PhD committee members Ilan Kremer and Chaim Fershtman, Stephan Lauermann and Sven Rady who were my hosts at Bonn Econ department where part of the research was conducted, members of Lauermann's lab for many kind and thoughtful suggestions, participants at the Hebrew U Econ PhD Seminar, at Mark Armstrong's search mini-course at Bonn and at the Econ Theory Summer School on Finance, as well as Danny Ben Shahar, Nathan Dortman, Darrel Duffie, Israel Feldman, Daniel Fershtman, Arthur Fishman, Holger Gerhardt, Alex Gershkov, Sergiu Hart, Carl Heese, Ziv Hellman, Dominik Karos, Deniz Kattwinkel, Dan Levin, Saul Lach, Avi Lichtig, Helene Mass, Jeffrey Mensch, Igal Milchtaich, Benny Moldovanu, Zvi Ne'eman, Assaf Patir, Dotan Persitz, Tobias Rachidi, Jose Scheinkman, Moses Shayo, Ofer Setty, Dror Speiser, Ran Spiegler, Eyal Winter, Ran Weksler, Asher Wolinsky, Nir Yosef and Boaz Zik. It has also benefited from financial support from the Hebrew University Economics Department, Ben Porath Foundation and Alrov Institute for Real Estate Research. Funding by the German Research Foundation (DFG) through CRC TR 224 (Project B04) is gratefully acknowledged.

†PhD Candidate in Economics, Hebrew University of Jerusalem. Contact: nikita.kotsenko@mail.huji.ac.il

1 Introduction

Uncertainty about market conditions prevails in many decentralized markets. Prime examples are labor and housing markets, where transaction data is not always available and where very similar transactions are infrequent. Participants must cope with this uncertainty by using public information and their own experience. This experience is often limited because most people search for a house or a job only a few times in their life. The interaction of public information and learning from experience cannot be captured by standard search models which assume market conditions are known in advance.

To study how public information affects outcomes, I develop a search, matching and strategic bargaining model in the tradition of Rubinstein and Wolinsky (1985). Buyers and sellers flow into a market at constant rates, where they search for partners and then strategically bargain over the division of the surplus. In their model, this entry rate is known, which is where my model differs. I assume the entry rate of buyers into the market is unknown. This implies that the equilibrium stocks of buyers and sellers and the probabilities of finding trade partners are also uncertain. First, I characterize the optimal behavior of fully rational agents under perfect Bayesian equilibrium and establish the existence and uniqueness of a full trade equilibrium. Second, I explore how public information affects equilibrium welfare, price dispersion, search duration and the path of reservation prices during search.

All previous studies of a decentralized market with aggregate uncertainty and learning assume at least one of the following: 1) At most one side of the market has any uncertainty (e.g., Duffie et al, 2017); 2) When two agents meet, they reveal their information to each other before agreeing on terms (e.g., Shneyerov and Wong, 2020); 3) Only two states are possible, uncertainty is about which of the states is true (e.g., Lauermann et al, 2017). None of these assumptions is innocuous. In my setting, both sides face the same kind of uncertainty. They do not exchange information but learn only from their outcomes. I place only weak restrictions on the prior distribution. In particular, I consider the class of prior distributions over the unknown parameter such that a more "precise" public signal induces agents to respond less to private signals.[1] A richer state space allows me to conduct comparative statics with respect to the precision of public information while keeping everything else fixed.[2]

These three features (two sided uncertainty, hidden continuation values and a rich state space)

[1] I define the notion of precision later on, but for the sake of intuition it can be regarded as the inverse of the variance of a public signal that is unbiased with respect to the true state.

[2] In a binary state space, reducing the prior variance must imply at least one of the following: either 1) the difference between the two possible states is reduced, which implies a change in the economic meaning of the two states and not just a change in beliefs, or 2) making one side more optimistic and the other side less optimistic from the onset. In either case, the effect of public signal precision cannot be studied in isolation.

have not been considered before in the same model and they generate new insights about an important class of economies.

Regarding welfare, I provide novel conditions such that welfare is increasing or decreasing with public information.[3] The welfare improving effect is related to the phenomenon of withdrawals in the housing market (Anenberg and Laufer, 2017) and of discouraged job seekers in the labor market (Blundell et al, 1998). These phenomenon are difficult to rationalize in standard search models where behavior is stationary. The present setting provides both a simple explanation and a potential remedy. As agents search, those that do not find a match for several consecutive periods conclude that conditions are worse than they initially expected. This pessimism, coupled with costly search, eventually causes some of them to quit the market. This occurs even when true market conditions are such that continued search is always optimal. More precise public information causes them to rationally be more doubtful about their private signals and thus postpones quitting. The negative effect is related to a congestion externality from increased participation (Ljungqvist and Sargent, 2007). Therefore, this study provides new insights about important phenomena in major markets.

Most of the price dispersion literature focuses on one of two distinct kinds of price dispersion: dispersion among posted prices (Varian, 1980; Pennerstorfer et al, 2020) and within the bid-ask spread (Glosten and Milgrom, 1985; Jovanovic and Menkvald, 2020). I focus on a third kind of price dispersion, which is dispersion within the bargaining spread, between the buyer's and the seller's positions (Kagel and Roth, 2016). I demonstrate it is increasing in public information under quite general conditions. This comes about because it is optimal for all agents to make offers that the most optimistic type on the other side of the market would accept (a semi-pooling equilibrium). This most optimistic type is more optimistic than the public signal, so increased public information makes this type **less optimistic** about gains from future search. This is the first effect. On the other hand, sequential rationality implies participants internalize their future beliefs into current continuation values. This implies that the future belief path and how it is affected by public information propagate back into the bargaining decisions of entering agents. This effect acts to mitigate the first effect because public information **improves future beliefs**. I provide general conditions such that the first effect dominates so public information reduces the continuation values of the most optimistic types. This causes the sellers of this type to accept lower prices and buyers to be willing to pay more, so the gap between

[3] Previous studies have identified several contexts in which public information can reduce welfare: it can reduce the scope for risk sharing in the economy (Hirshleifer, 1971; Eckwert and Zilcha, 2003); it can cause excessive coordination (Morris and Shin, 2002); it can crowd out private information in the sense of making individual actions less responsive to private information and therefore less informative (Amador and Weill, 2010).

the two prices increases. I use the one shot ultimatum bargaining protocol that allows this gap to be expressed in transaction prices. On the other hand, pooling of offers implies that the dispersion of beliefs among buyers and among sellers are not expressed in equilibrium transaction prices. Therefore, public information causes price dispersion to increase.

This result stands in sharp contrast to existing findings of public information reducing the bid-ask spread (Glosten and Milgrom, 1985; Stoll, 2000). However, I demonstrate that both phenomena can occur at the same time. To do so, I study a variant of my model with market makers as modeled by Duffie et al (2005). Market makers meet participants at a known rate and make an ultimatum offer which equals the reservation price of just-entered agents. These agents do not have any signals, so public information unambiguously increases their continuation values by improving their future belief path. Therefore, the bid-ask spread is decreasing with public information. Note however that in this case, price variance can still increase if market makers account for a small enough fraction of trade. I also show that public information can reduce price dispersion when price dispersion only reflects belief dispersion. A simple case occurs when only one side of the market makes offers, and when offers must equal reservation prices (as in Glaeser and Nathanson).

The above discussion of pooling offers also makes clear why prices are uninformative about market conditions, though they are informative about the extent of uncertainty agents face. On the other hand, the distribution of search durations is informative about market conditions and more importantly, it is informative about the likely effect of public information on welfare. If more precise public information increases the maximal search duration on at least one side but does not increase the average search duration on either side, it must also improve welfare. The intuition is that an increased maximal duration indicates postponed quitting, which increases the private utility of marginal agents, while the non-increased average search duration indicates that the added population does not cause a negative externality.

This study also has policy implications. First, it directly informs existing policy debates in the realms of public provision of information and market transparency. Statistical agencies, central banks, industry associations and market organizers all make decisions about which data to provide (see Duffie et al 2017 for a detailed discussion in the context of OTC markets). The welfare implications and diagnostic tools outlined in this study should be taken into account in such policy decisions. Second, this study raises the possibility that declining reservation wages during unemployment, a widespread phenomenon which harms both the unemployed and aggregate economic activity, can be mitigated by a specially designed informational intervention.

Pavan and Vives (2015) in surveying the literature on Information, Coordination and Market Frictions, point to dynamic problems of information acquisition as a key area for future research where "a lot remains to be done". They note that "we are still missing a tractable model" of dynamic learning and suggest that "extending the analysis to richer information ... structures" is "essential for a deeper understanding of the positive and normative properties of these economies." The present work contributes to this goal by developing a tractable model of dynamic learning with very weak restrictions on the prior distribution and provide positive and normative implications for an important class of economies. It is likely that the paucity of research in this area stems in part from the perception that such research projects must be difficult to comprehend and undertake. This study demonstrates that a model of dynamic learning can be quite straightforward in methods, arguments and implications.

The rest of this study is organized as follows: Section 2 sets up the model; Section 3 reviews the literature and highlights the contribution of this study; Section 4 describes the solution concept; Section 5 proves the comparative static results; Section 6 concludes.

2 Literature

2.1 Search, Matching and Bargaining

The relevant literature begins with Rubinstein and Wolinsky (1985) who introduce strategic bargaining into a search and matching market. Their main concern, which remains the main concern of the subsequent literature, is whether prices converge to the Walrasian prices when frictions vanish (patience $\delta \to 1$). Walrasian prices are 1 in a sellers' market and 0 in a buyers' market. They find that as impatience goes to zero, the price converges to a number between 0 and 1 that is determined by the matching chances of the two sides. So prices are not Walrasian unless the ratio of short to long side matching chances goes to zero.[4]

Wolinsky (1990) is the first to introduce uncertainty about the true state of the world to such a setting. He assumes the value of the traded item is known only to some of the agents. He asks if the market is able to aggregate information so that the uninformed end up trading at a price consistent with the true state. He assumes that only two states are possible (high and low) and only two bargaining

[4]It should be noted that their conclusions are disputed by Gale (1986) who shows that a frictionless market with bilateral strategic bargaining implements the Walrasian allocation. The differences between the models are too subtle to discuss here. See also Gale (2000).

positions are allowed (high and low). When $\delta \to 1$, uninformed agents can spend a lot of time insisting on the better deal and learning about the state from the bargaining positions of their counterparties. However, this makes uninformed agents relatively more common because informed agents trade quickly. This in turn implies that the uninformed are very likely to meet other uninformed which reduces the information value of each meeting. Thus, the value of acquiring information does not become negligible and some agents trade at the wrong price.

Lauermann et al (2017) revisit Wolinsky's question in a richer framework. Each seller may meet many buyers at the same time and no restrictions are placed on individual bids. Sellers hold first price sealed bid auctions with a secret reserve price as in Satterthwaite and Shneyerov (2008). Lauermann et al assume the item's value is known but the entry rate of buyers is not. The high and low states now refer to whether the entry rate of buyers is above or below that of sellers. Buyers learn from losing an auction with a certain bid while sellers learn from no being able to sell at a given reserve price. Importantly, the friction is exogenous exit so that as it goes to zero ($\delta \to 1$) the ratio the stocks of agents on the long side to the short side grows arbitrarily large. Price are Walrasian in this case.

In independent work conducted in parallel, Shneyerov and Wong (2020) also extend Rubinstein and Wolinsky (1985) but with the following differences compared to the present setting: time is continuous, agents have a monetary discount rate (in addition to exogenous exit), a specific parametric aggregate matching function is assumed (my results do not rely on such an assumption), the state space is binary (high or low), and each encounter starts with mutual revelation of search histories. The latter two are the most crucial. The binary state space implies that a reduced prior dispersion must change the economic meaning of the high and low states, so there is no scope for a pure information effect. This is because when only two states are possible, a reduced prior dispersion must imply that the two possible states are more similar to one another in economic terms, and also closer to a balanced market, where the surplus is divided equally.

Mutual revelation of information ensures the existence of a full trade equilibrium for arbitrarily low levels of friction $(1 - \delta + r)$, which is their main focus. The goal of Shneyerov and Wong is to study convergence to Walrasian prices as frictions vanish in the presence of uncertainty about buyer entry rates. In contrast, my focus is on welfare and price dispersion in the presence of frictions. Welfare is trivially increasing in their model as more agents end up trading due to lower exogenous exit. They show that there is only one full trade equilibrium but do not provide conditions such that full trade is the unique equilibrium in some larger class, as I provide here.

Full trade is important because models of this kind face a distinct tractability challenge. Optimal behavior depends on each type's beliefs (subjective distribution) about the distribution of types, while the distribution of types depends on the behavior of each type. Therefore, each type's posterior can be a function of distributions of distributions (of distributions...). In addition, posteriors must include both binary signals (matching or not matching) and discrete but non binary signals (offers and responses which signal the distribution of types on the other side). Therefore, with finite histories, such posteriors do not fall into any known distribution families and therefore are difficult to study analytically. Full trade, along with one shot bargaining schemes and anonymity, ensures that signals are binary (meeting or no meeting), thus simplifying posteriors. The binary state space and the mutual revelation of information also contributes to making equilibrium prices reasonably tractable.[5]

I make several contributions to the literature. First, I analyze the welfare implications of public information in decentralized markets with aggregate uncertainty. In so doing, I make a novel connection between such markets and the empirical phenomena of withdrawals in the housing market and discouraged workers in the labor market. Second, I demonstrate that conclusions can be reached without little or no restrictions on the state space. Third, I provide more general conditions for equilibrium uniqueness than Shneyerov and Wong.

2.2 Public Information and Welfare

In terms of the comparative statics I focus on, a few other strands of the literature are related. First is the literature on the welfare effects of public information. Most of this literature focuses on static models which are less relevant to the current study (see Pavan and Vives, 2015 for a review). The most closely related studies to my own are Duffie et al (2009 and 2017). Duffie et al (2009) study the welfare effects of public information in a dynamic setting with search and matching. However, their setting does not involve trade between meeting agents. Instead, agents mutually reveal their information upon meeting, and this information is the source of private signals. Based on their information, agents choose optimal search intensities and a terminal action. A related study by Duffie et al (2014) has bilateral trade but not public information. Duffie et al (2017) study the welfare effects of public information on a dynamic exchange economy with search and matching and provide conditions for information to be welfare improving. However, their model differs in several key respects. First, public information

[5] Approaches not including full trade equilibria exist: Wolinsky (1990) restricts possible bargaining stances to be high or low; Majumdar et al (2016) assume point beliefs; Lauermann et al (2017) focus on the limit where $\delta \to 1$ which facilitates rapid learning as the ratio of long side to short side populations and goes to infinity. All three impose a binary state space.

is either not provided or is fully revealing. Second, the information concerns the seller's costs so only one side has any meaningful uncertainty.

I contribute to this literature by demonstrating a new channel from public information to welfare which is relevant to major markets such as housing and labor.

2.3 Information and Price Dispersion

This study is also relevant to the literature on information and price dispersion in bilateral trade. Most of this literature, from Stigler (1961) to Pennerstorfer (2020), focuses on one sided search (see Baye et al (2004) for a review) where sellers post prices and buyers search among them. The typical finding is that price dispersion has an inverted U shape with respect to information (first reported by Diamond, 1971). Glosten and Milgrom (1985) study a model with informed traders, uninformed traders and uninformed market makers. Their setting does not have search and matching frictions and the unknown is the value of the asset. They assume that public information reduces the information advantage of informed traders over the market markets and find that it reduces the bid-ask spread. In my setting public information reduces the dispersion of beliefs but this does not have to result in reduced price dispersion. This is in part because price dispersion at the bargaining "spread" between the buyer and the seller is distinct from price dispersion at the bid-ask spread. See more discussion of this issue after the price dispersion comparative static in subsection 5.1.4. Duffie et al (2005) study a search and matching market with buyers, sellers and market markers, but their setting has no uncertainty and so no role for public information. I use their formulation of market markers in an extension to show the effect of public information on the bid-ask spread. This allows me to show that the bid-ask spread and the bargaining spread can be moved by public information in different directions (see subsection 5.1.5).

A new results of this study is to provide sufficient conditions such that price dispersion is monotonically increasing with public information.

2.4 Declining Reservation Prices

Finally, this study is also relevant to the literature on declining reservation prices during search. This phenomenon has been observed in the field and in laboratory settings (see Brown et al 2011 for a review of early evidence and for new laboratory findings). It is important in labor markets where declining reservation wages during unemployment can lead to lower subsequent wages as well as to lower participation due by discouraged workers. Gronau (1971) rationalizes declining reservation prices as caused by a finite search horizon. Burdett and Vishwanath (1988) show the same phenomenon can be rationalized by learning even if the horizon is infinite. Their model assumes bids are distributed normal and are exogenous. Although they do not mention it, one can easily show that higher prior precision mitigates declining reservation prices in their model. This can have important policy implications. Brown et al (2011) provide suggestive evidence from the lab and the field that uncertainty and learning are implicated in declining reservation prices.

A new results of this study is to demonstrate that reservation prices decline during search on both sides of the market and regardless of the particular signal structure. Rather, it is a direct consequence of sequential learning. Furthermore, it can be mitigated by providing more accurate public information. To the best of my knowledge, this point has not been made before.

3 Setup

Time is discrete and the market is perpetual. At the start of time, nature draws a single realization x from the random variable X. This realization determines the "state" (the meaning of which will be explained shortly) and remains fixed forever. While X is common knowledge, x is not observed by anyone.

The market is populated by buyers and sellers with linear utilities. Buyers (B) and sellers (S) have potentially infinite horizons, but fraction $1-\delta_j \in (0,1), j = \{B,S\}$ of them exit exogenously each period. Buyers have unit demand for an indivisible item and value the object at 1 (it may be convenient to think of the item as a house). Sellers have unit supply of the item at cost 0. Each period, a new mass $N_{S,0}$ of sellers and a new mass $N_{B,0}(x)$ of buyers enter the market to join existing buyers and sellers. The subscript 0 indicates that the fixed entry rate is always equal to the population of side j who had 0 matching opportunities. Because all the analysis is in steady state, I ignore calendar time from the onset.

Both entry rates are fixed, but only $N_{S,0}$ is known. The function $N_{B,0}(x)$ is known and increasing in the realized state x, so that information about x provides information about $N_{B,0}(x)$. As will be made clear shortly, knowing x implies knowing $N_{B,0}(x)$, which enables one to calculate the equilibrium matching rates and population stocks.

3.1 Order of events

The order of events in each period is as follows:

1. $N_{S,0}$ sellers and $N_{B,0}(x)$ buyers enter and join the agents who remain from previous periods. All agents are identical except for their role of buyer or seller and their histories which cause different beliefs. In what follow, I refer to the entire stocks of agents jointly.

2. Each agent chooses whether to pay c_j and continue search or to not pay and quit (this and subsequent events apply to new and old agents alike).

3. Sellers and buyers are randomly matched into pairs at the individual per period rates $\omega_j(x) \in (0,1)$.[6]

4. Matched pairs bargain with private histories. They flip a fair coin to make a take it or leave it offer. If the receiver agrees, both consume and exit. If not, the meeting dissolves and agents re-join the market.

5. A fraction $1-\delta_j \in (0,1), j=\{B,S\}$ of the remaining agents exit exogenously.[7]

3.2 Information

Generally speaking, each agent only observes what happens to her, so she knows when she enters, matches, receives an offer, gets an exit shock and so on. She does not observe what happens to others and what others decide, other than during bargaining where bids and responses are observed by both parties, though not by outside parties. In addition, agents understand the model perfectly and know all parameters except for the realization of x. I am assuming throughout that agents use Bayes' Rule to update their beliefs but other updating rules may be equally applicable.

The common prior over x has a PDF $f_Y(x|X,\pi(k))$ that depends on the true state distribution X and on a public signal $\pi \in (0,1)$, with $E[\pi]=x$. The distribution of π depends on a parameter k.

Define the posterior mean belief about the per period matching probability as follows:

[6] As usual, $\omega_j(x) \equiv \frac{M(N_S(x),N_B(x))}{N_j(x)}$, where $N_j(x)=\sum_{t=0}^{T_j-1} N_{j,t}(x)$ is the total mass of agents on side j, $T_j \leq \infty$ is the endogenous quitting period, and $M(,)$ is the matching function that determines the mass of matches each period based on the total populations of buyers and sellers. For now, I make no assumptions about $M(,)$ except that it is fixed (conditional on populations), weakly increasing in each argument and is known to agents. Also, matching is random in the sense that each pair that contains one buyer and one seller is as likely as any other such pair.

[7] This threat of termination induces impatience. Brown et al (2011) verify experimentally that subjects respond very similarly to monetary discounting as to termination-threat discounting.

$\mu_{j,h} \equiv \int_0^1 \omega_j(x) f_Y(x|X, \pi(k), h)dx$, where $h = (s_1, s_2, ...)$ is some history of private signals s_i observed during participation (matching or not matching, bids received and rejected).

I will often refer to k as the precision of the prior for a reason that will be clear shortly. However, k can be any parameter affecting π such that the following assumptions hold:

A1 For any two levels of k, (k_0, k_1) such that $k_0 < k_1 < \infty$, and for any two histories, h_0 and $h_1 = (h_0, s)$, where s is a private signal, the effect of the private signal s has a weaker effect on the posterior mean under k_1 relative to k_0:[8] $|\mu_{j,h_0}(k_0) - \mu_{j,h_1}(k_0)| > |\mu_{j,h_0}(k_1) - \mu_{j,h_1}(k_1)|$

A2 X, π are such that whenever $c_j > 0$, there exists a history h of binary signals such that $\mu_{j,h} < c_j$ for $j = \{B, S\}$. This is analogous to assuming that the state space and the public signal assign a positive probability to a non empty set of states where sellers have arbitrarily low beliefs and to another non empty set of states where buyers have arbitrarily low beliefs.

A1 and A2 together allow me to abstract from assuming a particular joint distribution for the information structure.

As an example of A1, one can consider a (post-public signal) prior of $f_Y(x|X, \pi(k)) \sim Beta(yk, zk)$ with k not too small (The Beta distribution is commonly used to model priors over beliefs). In such a case, one can verify that increasing k reduces the difference in posterior mean between two agents on the same side whose histories differ by at most one binary signal. A similar result can be readily obtained when the signals are drawn from the normal distribution with unknown mean that is also drawn from a normal distribution with precision k. For this reason, the reader may find it convenient to think of k as the precision of the common prior.

3.3 Discussion

As a concrete example for the setting, consider a homeowner who would like to sell her home. To decide on a bargaining stance, she must consider the rate of arrival of buyers and the outside option of those buyers. To do so, she consults a publicly available data source, which is the realized public signal π_0, knowing that it is also available to all other agents. In addition, she uses her knowledge of the setting and her ability to acquire additional private signals during search.

It is reasonable to suppose that in the relatively short time she will spend on the market, the

[8] For the results in this study apart from the equilibrium uniqueness result it is enough to assume this holds for binary signals ($s_i \in \{0, 1\}$).

informativeness of the available information will not change in any meaningful way. It is also possible that sellers do not consult public data more than once during their search as such consultation is costly in time and effort. Therefore, it is reasonable to assume that π_0 is fixed.

Also note that housing search is segmented in the sense that a particular buyer does not examine all possible homes but focuses her search on houses with similar characteristics (Piazzesi et al 2020). It is therefore a useful simplification to suppose all buyers are identical in their preferences, in the rate at which they find other sellers and that those sellers offer essentially identical houses. This simplification is especially useful because precise information on the distribution of buyer preferences and house unobserved quality is not publicly available.

4 Equilibrium

I focus on a steady state full trade perfect Bayesian equilibrium, which I will call FTE. The focus on a full trade equilibrium provides a clear intuition for novel and important economic results.

In contrast, a partial trade equilibrium is complicated by the fact that posteriors must include both binary signals (matching or not matching) and discrete but non binary signals (offers and responses which signal the distribution of types on the other side). Therefore, with finite histories, such posteriors do not fall into any known distribution families and therefore are difficult to study analytically. I follow Shneyerov and Wong (2020) in focusing on a full trade equilibrium. See Lauermann et al (2017) for an ingenious effort in establishing some partial trade results in a similar context, though at the cost of assuming a binary state space. While interesting, their results are not a substitute for the present work because they concern the effects of $\delta \to 1$ and not of a more precise prior, which their setting does not accommodate. I discuss candidate partial trade equilibria in Appendix E. Though largely conjectural, the discussion does imply that the main insights of the analysis in full trade extend to some partial trade equilibria.

4.1 Definition: FTE

A participant's type is fully characterized by her role of buyer or seller ($j = \{B, S\}$) and her history h which includes the common prior (X), the public signal (π), the number of periods in the market ($t = \{0, 1, ...\}$), the number of meetings ($m = \{0, 1, ..., t\}$), the bids received and the bids made in each meeting.[9]

A perfect Bayesian steady state full trade equilibrium in cut-off strategies consists of:

1. **FTE Populations:** function $N_{j,h}(x)$ from type to steady state mass of agents of this type that are in the market conditional on the parameters and the realized state, that are consistent with FTE strategies.[10]

2. **FTE Beliefs:** Posterior probability density functions $f_Y(x|h)$ that are consistent with Bayes' rule, with FTE strategies and with FTE populations;

3. **FTE Strategies:** Offer and reservation strategies $o_{j,h}, r_{j,h}$ that are consistent with FTE beliefs such that

   (a) no unilateral deviation is strictly welfare improving given beliefs and

   (b) offers consistent with FTE play are always accepted.[11]

Proposition 1 (FTE): There exists k^* such that for any $k \geq k^*$, a PBE in cutoff strategies exists and any such PBE has full trade.

A full treatment is given in the Appendix E. Here I give an intuitive description. The rest of the analysis will be assuming FTE play.

4.2 Intuition for Existence and Uniqueness

The essence of the arguments for both existence and uniqueness relies on a trade-off between making an offer that is always accepted (safe) or making an offer that is sometimes rejected but yields more utility if it is accepted (risky). The risky offer is better when the difference between reservation prices on the other side is large (heterogeneity). The more similar are the agents on the other side, the lower is the incentive to make an offer that some of them will reject. The safe offer is better when it is more likely that one is facing someone who will reject. A higher k causes lower heterogeneity and by

[9] I am counting tenure in the market as the number of times an agent was present when meetings were announced. So agents enter the market with tenure 0 but end their first period in the market (if they reach the end without matching) with tenure 1.

[10] I follow Shneyerov and Wong (2020) in not imposing steady state conditions as part of the equilibrium definition. See Appendix C for a detailed proof that steady state always arises in this market as a consequence of FTE play.

[11] If an agent is asked to respond to an off-path offer (an offer that is not consistent with FTE play), she responds according to her reservation price. The reservation price is calculated as usual, conditional on her tenure and on meeting in this period (conditioning on receiving an off path offer and on the size of the offer does not change the posterior because it is a zero probability event).

choosing an appropriate k one can make heterogeneity arbitrarily low.[12] On the other hand, k does not have the same effect on the probability of rejection. The proportion of agents who are the most optimistic about their future chances is always strictly positive as long as $\delta < 1$, which I am assuming throughout. So one can always find a lower bound for k such that for any k above this bound, the gain from making the risky bid is not enough to cover the probability of rejection. Note that all the above does not rely in any way on an assumption that everyone else plays according to FTE. Rather, it implies that FTE play can be made a dominant strategy by a suitable choice of k. For such a choice, FTE exists and is unique.

For a graphical illustration see Figure 1 below. It illustrates the trade-off faced by a buyer in FTE when she is selected to make the offer. She is facing a seller with unknown tenure and it is the first meeting for both of them. The buyer's utility if her offer is accepted is on the vertical axis and the probability her offer is accepted is on the horizontal axis. The red dots indicate two possible choices: the risky offer $r_{S,2}$ which yields utility $1-r_{S,2}$ with probability of acceptance below 1 and the safe offer $r_{S,1}$ which yields the lower utility $1 - r_{S,1}$ with probability of acceptance 1.[13] A higher k reduces the vertical distance by making sellers less heterogeneous. If we zoom out far enough so that many tenures are included, the black step function would appear smooth, with a negative slope left of probability 1 and then vertical as it is now. A higher k reduces the average "slope" of the step function left of probability 1. If k is very large, the trade-off step function will be very similar to a horizontal line. Then, a very small reduction in the offer yields a very high reduction in the probability that the offer is accepted. This naturally leads to the safe offer being made.

A similar logic applies to ruling out partial trade equilibria. To see this, image that $r_{S,1}$ now denotes the reservation price of the most optimistic seller, while $r_{S,2}$ denotes that of the second most optimistic seller and so on. The rest of the intuition is as above.

4.3 Intuition for Steady State

When writing $M \equiv M(N_S, N_B)$ for the mass of period period matches, the condition for the steady state total stock on side j is:

$$N_{j,0} = M + N_{j,T_j} + (N_j - M)(1 - \delta_j)$$

[12]As long as $k < \infty$, k is never too high for any of the arguments in this study to be correct. At the same time, k can often be quite small in practice. See Appendix E for more details.

[13]$r_{S,t}$ denotes the reservation price of a tenure t seller in their first match. I also sometimes write this as $r_{S,t,1}$ when confusion is possible. Note that in FTE a type is fully captured by the tuple (j, t, m).

Figure 1: Optimal Offer Trade-off

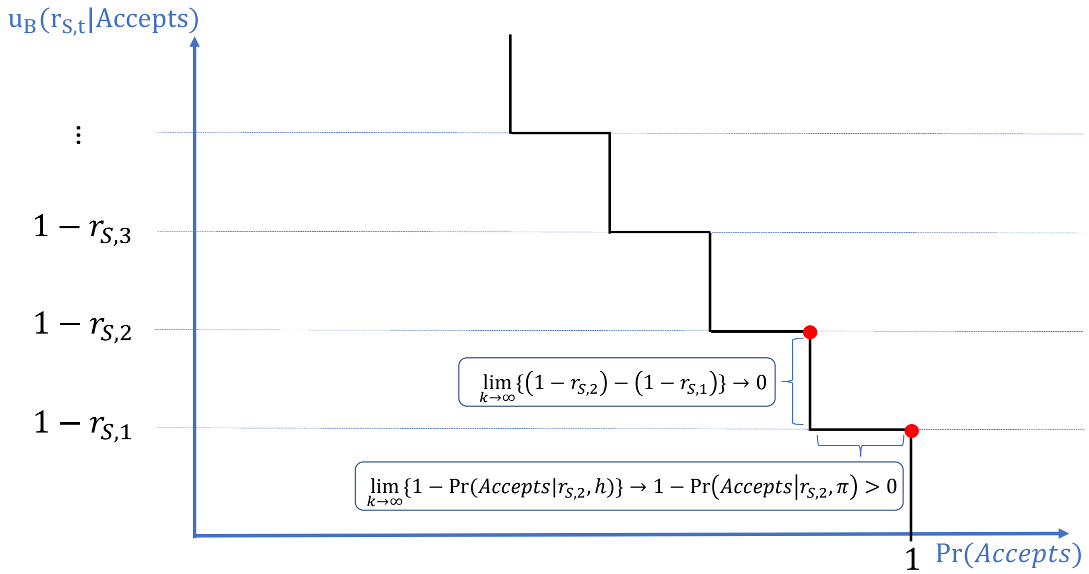

Notes: the horizontal axis is the subjective probability that an offer is accepted by a seller, from the perspective of the buyer; the vertical axis is the utility to the buyer given that her offer is accepted. In FTE this is the first meeting for both so the third subscript (m) is omitted. $Pr(Accepts|r_{S,2}, h)$ denotes the probability a buyer with history h assigns to offer $r_{S,2}$ being accepted. $Pr(Accepts|r_{S,2}, \pi)$ denotes the same except that now beliefs are conditional on the public signal being almost surely equal to the true state x, which also makes the history irrelevant.

A fixed mass of agents $N_{j,0}$ is added every period. Assuming FTE play, a fixed fraction of the stock is subtracted every period. The subtracted mass includes those that chose to quit (N_{j,T_j}), those that matched (M), and those that did not match but exited exogenously $(N_j - M)(1 - \delta_j)$. Recall that $N_j \equiv \sum_{t=0}^{T_j-1} N_{j,t}$.

To see why convergence to steady state occurs, note that if $(N_j - M)$ is above (below) its steady state level, then $(N_j - M)(1 - \delta_j)$ is also above (below) its steady state level so that next period $(N_j - M)$ is smaller (larger). This pushes population towards the steady state level. Therefore, for any initial population stocks, the long term stock of agents is the sum of a convergent geometric series.

Note that the effect of M on the population stock does not negate this process because if N_j is above (below) its steady state level then M is weakly above (below) its steady state level as well because of the standard assumption that $M(N_S, N_B)$ is weakly increasing in each argument. Because entry rates are in masses, the number of agents on each side is always infinite. This ensures that the law of large numbers applies, so that the fraction of agents who experiences a certain event always equals the probability of a given agent to experience this event.

Consequently, the long run total stocks are fixed, so that type stocks are fixed as well, as they depend only on total stocks and model parameters. Thus the market converges to a steady state where the population of each type is constant. I focus on this limit.[14]

4.4 Remarks

Note that in FTE, no agents continues to search after a meeting. Therefore, every type on the equilibrium path can be described by the tuple (j, t, m).

Also note that because bargaining is conducted with private histories, the only offer strategy that supports an FTE is where all sellers, regardless of their history, make an offer that leaves the most optimistic buyer indifferent (as customary, I am assuming agents accept if they are indifferent). Buyers likewise make an offer that leaves the most optimistic seller indifferent. If sellers ask for more, the most optimistic buyers will reject, so we are not in FTE. If they ask for less, they leave money on the table. The most optimistic agent in FTE after meetings are announced is a tenure 1 agent with 1 meeting. This implies that all equilibrium offers are $o_{j,t,1} = r_{-j,1,1}$.

Therefore, a seller's expected utility from a meeting is given by:

[14] See Appendix C for a formal and detailed proof.

$$u_S = \frac{1}{2} r_{B,1,1} + \frac{1}{2} r_{S,1,1}$$

This follows because with probability $\frac{1}{2}$ the seller is chosen to make the offer, in which case she offers to be paid $r_{B,1,1}$ and the buyer accepts, and with probability $\frac{1}{2}$ the buyer makes the offer, in which case she offers to pay $r_{S,1,1}$ and the seller will agree. Because the seller provides the item at cost 0 and she is risk neutral, her utility is the average between $r_{B,1,1}$ and $r_{S,1,1}$.

The buyer's case is similar except for the fact that the buyer values the item at 1 and so, her expected utility from a match is:

$$u_B = \frac{1}{2}(1 - r_{B,1,1}) + \frac{1}{2}(1 - r_{S,1,1})$$

4.5 Value Functions

In the present setting, an agent makes decisions at two points in time: before paying the search cost and if she is in a meeting. I now derive the value function for these two cases.

4.5.1 In a meeting

Now I derive the FTE continuation value of an agent on side j, after meeting in her t'th period on the market in the event that bargaining breaks down:

$$V_{j,t,1} \equiv (1 - \delta_j) \cdot 0 + \delta_j \cdot (-c_j + \mu_{j,t,1} \cdot u_j + (1 - \mu_{j,t,1}) \cdot V_{j,t+1,1}) \tag{1}$$

To understand equation (1), first note that with probability $1 - \delta_j$ the agent exits at the end of period t with utility 0. With probability δ_j she continues to period $t+1$, where she must first pay c_j to not quit with utility 0. If she does so, she has subjective probability $\mu_{j,t,1}$ to meet a trading partner. If a meeting occurs, her expected utility from the meeting is u_j and then she exits. A meeting does not occur with subjective probability $\mu_{j,t,1}$, in which case her subjective continuation utility is $V_{j,t+1,1}$.[15]

A more concise formulation is:

[15] Note that $V_{j,t+1,1}$ already takes into account the information of not meeting in period $t+1$. This is necessary for sequential rationality as prescribed by a Perfect Bayesian equilibrium. Recall that sequential rationality requires that the agent's strategy in each node be an optimal response given the information that she has at this node. Therefore, when formulating $V_{j,t,1}$, the agent must take into account the fact that her information set, if she happens to reach $t+2$ will contain an additional signal, which on the equilibrium path, will be that she did not have a meeting.

$$V_{j,t,1} \equiv d_{j,t,1} \cdot u_j - g_{j,t,1} \cdot c_j \tag{2}$$

Where $d_{j,t,1}$ is the subjective probability to eventually meet a trading partner and $g_{j,t,1}$ is the expected number of periods until exiting the market for any reason.

To arrive at (2), I repeatedly substitute the appropriate RHS for $V_{j,t+1}, V_{j,t+2}, ...$ in (1) to get (2) , where

$d_{j,t,1} \equiv \delta_j \mu_{j,t,1} + \sum_{\tau=2}^{T_j-1} \delta^\tau \Pi_{i=t}^{t+\tau-2}(1-\mu_{j,i,1}) \cdot \mu_{j,t+\tau-1,1}$

$g_{j,t,1} \equiv \delta_j + \sum_{\tau=2}^{T_j-1} \delta_j^\tau \cdot \Pi_{i=t}^{t+\tau-2}(1-\mu_{j,i,1})$

T_j is the endogenous terminal period in which an agent chooses to quit the market.

$d_{j,t,1}$, has an intuitive form. δ_j is the probability to not exit exogenously in a given period and $\mu_{j,t,1}$ is the expectation of the per period meeting rate after one meeting in t periods. So $d_{j,t,1}$ equals the probability to not exit exogenously and to meet in period $t+1$ (according to beliefs based on t signals) or to not meet in period $t+1$ (this is the term in the product) and to not exit exogenously and meet in period $t+2$ and so forth. The crucial points are that 1) $d_{j,t,1}$ depends on the future path of beliefs and 2) under FTE $d_{j,t,1}$ is exogenously determined by the parameters and the information structure.

$g_{j,t,1}$ is analogous and equals the expected number of periods the participant expects to pay c_j before exiting for any reason. $T_j - 1$ the final period in which participants can be matched before deciding to quit instead of paying the search cost.

Remark:

1. Note that $d_{j,t,1}$ differs from $d_{j,t+1,1}$ only to the extent that $\mu_{j,\tau,1}$ differ from $\mu_{j,\tau+1,1}$ for all $\tau = \{t, t+1, ...\}$. The same is true for the extent that $g_{j,t,1}$ differs from $g_{j,t+1,1}$. Note also that $V_{j,t,1} - V_{j,t+1,1} = (d_{j,t,1} - d_{j,t+1,1}) \cdot u_j - (g_{j,t,1} - g_{j,t+1,1}) \cdot c_j$. So the difference between $V_{j,t,1}$ and $V_{j,t+1,1}$ also comes down to the difference between $\mu_{j,\tau,1}$ and $\mu_{j,\tau+1,1}$. Therefore, assuming a larger k reduces the difference between $\mu_{j,\tau,1}$ and $\mu_{j,\tau+1,1}$ (which is a special case of A1), it also reduces the differences between $V_{j,t,1}$ and $V_{j,t+1,1}$. Recall that k reducing the difference between continuation values is the key for the existence of a full trade equilibrium.

2. Note that $\mu_{j,t-1,m} > \mu_{j,t,m}, \forall t > 0, \forall m$: an agent's posterior mean matching probability is lower

after not matching in an additional period.

4.5.2 Before paying the search cost

Now suppose an agent did not match in period t and is now in period $t+1$ but before matching is announced. Therefore, this agent has type $(t, 0)$.

$$V^c_{j,t,0} \equiv -c_j + \mu_{j,t,0} \cdot u_j + (1 - \mu_{j,t,0}) \cdot V^c_{j,t+1,0} \tag{3}$$

$$= \delta^{-1}(d_{j,t,0} \cdot u_j - g_{j,t,0} \cdot c_j) \tag{4}$$

Where $d_{j,t,0}, g_{j,t,0}$ are defined analogously to the case inside the meeting. There are only two differences between $V^c_{j,t,0}$ and $V_{j,t,1}$. First, there is the absence of a successful matching attempt. Second, the agent has already survived the exogenous exit shock at the end of his period t in the market.

4.6 Weak Prior Mean Neutrality of k

I now introduce a further assumption that is used to simplify some of following proofs.

A3 Define π^* such that for $k_0 < k_1 < \infty$, $\mu_{j,0,0}(\pi^*, k_0) = \mu_{j,0,0}(\pi^*, k_1)$. I assume there exists a constant $\overline{\pi} > 0$ such that if $|\pi_0 - \pi^*| \leq \overline{\pi}$ then

1. $\mu_{j,t,0}(k_0) < \mu_{j,t,0}(k_1)$ for all t, j
2. $\mu_{j,1,1}(k_0) > \mu_{j,1,1}(k_1)$ for all j

To understand this assumption, note that the realized signal might not be prior neutral, in the sense that making the signal more informative, makes one side's prior mean higher and the other side's prior mean lower. In such a case, k has two effects on $\mu_{j,t,m}$: 1) through slower learning due to A1 and 2) due to the change in the prior. A3 is meant to specify conditions, such that the first effect dominates. Therefore, when all the signals are negative, μ is increasing in k and when the only signal is positive, μ is decreasing in k. This assumption holds for $f_Y(x|X, \pi(k)) \sim Beta(yk, zk)$.

Implication: Because $d_{j,t,0}$ is increasing and $g_{j,t,0}$ is decreasing in the sequence $\mu_{j,\tau+1,1}$ for $\tau = \{t, t+1, ...\}$, A3 implies that if $|\pi_0 - \pi^*| \leq \overline{\pi}$ then $V^c_{j,t,0}(k_0) < V^c_{j,t,0}(k_1)$.

4.7 A Summary of Intermediate Results

A few intermediate results are useful to keep in mind:

1. $o_{j,t,1} = r_{-j,1,1}$. The only offer strategy that supports equilibrium is one which targets the most optimistic agents on the other side.

2. $\mu_{j,t,m} > \mu_{j,t+1,m}$: An agent's posterior mean matching probability is lower after not matching in an additional period. This implies that the belief path of an agent is declining (sampled at the same stage of consecutive periods).

3. If $c_j > 0$, agents with enough negative signals will quit. Denote this threshold

$$T_j = min\{t|V^c_{j,t,0} < 0\}$$

.[16]

4. If $|\pi_0 - \pi^*| \leq \overline{\pi}$ then $V^c_{j,t,0}(k_0) < V^c_{j,t,0}(k_1)$

5. If $|\pi_0 - \pi^*| \leq \overline{\pi}$ and $\mu_{j,t,m}$ is differentiable with respect to k then $\frac{\partial}{\partial k}\mu_{j,1,1} < 0$

5 Comparative Statics

5.1 Price Dispersion

Proposition 2 provides sufficient conditions for higher k causing higher price dispersion. After stating the proposition, I provide an intuitive explanation before stating the formal proof.

I define price dispersion as the difference between the two equilibrium transaction prices: $PD \equiv r_{B,1,1} - r_{S,1,1} > 0$. The fact that finding another partner is not guaranteed dictates that the offerer has some market power. This market power ensures that surplus is not split evenly in any given meeting which explains why $PD > 0$ (otherwise we would have $r_{B,1,1} = r_{S,1,1}$). See Appendix B for details.

5.1.1 Proposition 2

1. If for all $j = \{B, S\}$ and for all tenures $l = \{1, 2, ...\}$, $\frac{\partial}{\partial k}\mu_{j,l,1}$ exist and $\frac{\partial}{\partial k}\mu_{j,l,1} < \infty$, and

2. $\frac{\partial}{\partial k}\mu_{j,1,1} < 0$

[16]This is directly implied by A2

There exists (δ_S^*, δ_B^*) such that for all $\delta_S \in (0, \delta_S^*)$ and all $\delta_B \in (0, \delta_B^*)$, we have that:

$$\frac{\partial}{\partial k} PD > 0$$

5.1.2 Intuition

Note that agents with history (1,1) have observed one favorable signal in one attempt so that their beliefs are high relative to the public signal. Second, recall that a higher k weakens the effect of private signals, which acts to make agents with a history of (1,1) less optimistic when k is higher. Next, recall that when sellers are less optimistic, they are willing to accept lower prices so that $r_{S,1,1} \downarrow$ while buyers, when they are less optimistic, are willing to pay more $r_{B,1,1} \uparrow$. This is because optimism and pessimism here refer to prospects from continued search. When such prospects appear worse, this increases the incentive to trade quickly. When $r_{S,1,1} \downarrow$ and $r_{B,1,1} \uparrow$, the difference between the two prices increases so that price dispersion increases. I call this the "present" effect.

However, k has another effect on reservation prices because it makes future beliefs, conditional on being in the market, better. In PBE, agents take this into account. This causes $r_{S,1,1} \uparrow$ and $r_{B,1,1} \downarrow$ which counteracts the "present" effect. I call this the "future" effect. In principle, either effect could dominate. Which ends up dominating depends on the probability of being in the market long enough to enjoy the (perceived) better future prospects. This in turn depends on δ_j and $\mu_{j,t,1}$ for $t > 1$. If δ_j is not very high, exogenous exit is likely, so the future is not long enough and the "present" effect dominates. The proof below formalizes this idea. I note below that under further assumptions, one can show that the requirement that δ_j is not too large places only a mild restriction on the allowable range of δ_j.

Another approach to prove the same thing is to note that if $\mu_{j,t,1}$ is not too low, matching is likely, so that again the future is not long enough and the "present" effect dominates. I do not pursue this approach because it requires less transparent conditions.

5.1.3 Proof

In FTE price dispersion equals $r_{B,1,1} - r_{S,1,1}$, so that:

$$\frac{\partial}{\partial k} PD = -2\left[\frac{\frac{\partial}{\partial k} d_{S,1,1}(1 - d_{B,1,1})^2 + \frac{\partial}{\partial k} d_{B,1,1}(1 - d_{S,1,1})^2}{(2 - d_{B,1,1} - d_{S,1,1})^2}\right]$$

(see Appendix D for a derivation of the above expression. This proof assumes $c_S = c_B = 0$ for ease of notation. A proof with $c_j > 0$ is analogous.)

Note: $d_{j,1,1} \in (0,1)$ because it is the probability assigned by type $(j,1,1)$ to meeting before she exits exogenously or quits.

This implies that if $Sign[\frac{\partial}{\partial k} d_{S,1,1}] = Sign[\frac{\partial}{\partial k} d_{B,1,1}]$ then $\frac{\partial}{\partial k} PD = -Sign[\frac{\partial}{\partial k} d_{S,1,1}]$, because all other terms in $\frac{\partial}{\partial k} PD$ are positive.

So I focus on the signs of $\frac{\partial}{\partial k} d_{S,1,1}$, $\frac{\partial}{\partial k} d_{B,1,1}$. One can verify that:

$$\frac{\partial}{\partial k} d_{j,1,1} = [\delta_j \cdot \frac{\partial}{\partial k} \mu_{j,1,1} \cdot (1 - d_{j,2,1})] + \{\sum_{l=2}^{\infty} \delta_j^l \cdot \frac{\partial}{\partial k} \mu_{j,l,1} \cdot \Pi_{i=1}^{l-1} (1 - \mu_{j,i,1}) \cdot (1 - d_{j,l+1,1})\}$$

Note that when $\frac{\partial}{\partial k} \mu_{j,1,1} < 0$ as assumed, the term in [] is negative. Now note that $\Pi_{i=1}^{l-1}(1 - \mu_{j,i,1})$ and $(1 - d_{j,l+1,1})$ are bounded between (0,1) for all l and that δ has a minimal exponent of 2. Therefore, the term in {} goes to zero faster than the term in [].

Therefore, $\frac{\partial}{\partial k} \mu_{j,1,1} < 0$, $\exists \delta_j^* : \forall \delta_j \in (0, \delta_j^*)$, $\frac{\partial}{\partial k} d_{j,1,1} < 0$, which implies that $\exists (\delta_S^*, \delta_B^*) : \forall \delta_S \in (0, \delta_S^*), \forall \delta_B \in (0, \delta_B^*)$,

$$\frac{\partial}{\partial k} PD > 0$$

QED

If one assumes that $f_Y(x|X, \pi(k)) \sim Beta(yk, zk)$, one can verify that δ_j^* can be quite high and often above 0.99. The main reason is that $\Pi_{i=1}^{l-1}(1 - \mu_{j,i,1})$ turns out to be quite small. This is because it starts from $\mu_{j,1,1}$, which, for priors that are not too skewed, is quit large. For example, for the uniform prior, $\mu_{j,1,1} = \frac{2}{3}$. $\mu_{j,i,1}$ for higher i's decline rather slowly even for a dispersed prior like the uniform.

5.1.4 Conditions

Proposition 2 is stated in terms of endogenous quantities for ease of presentation and because this is the most general statement available. I now clarify what these conditions require.

Condition 1 has two components. One is the existence of a derivative of $\mu_{j,l,1}$ with respect to k and the second is that this derivative is bounded from above. Recalling that $\mu_{j,t,1} \equiv \int_0^1 \omega_j(x) f_Y(x|X, \pi(k), t, 1) dx$, the first part holds whenever $\frac{\partial}{\partial k} f_Y(x|X, \pi(k), t, 1)$ and $\frac{\partial}{\partial k} \omega_j(x)$ exist. $\frac{\partial}{\partial k} \omega_j(x)$ exists whenever $V_{j,T_j,0}^c < 0$ and $V_{j,T_j-1,0}^c > 0$ so that a small change in k does not affect quitting decisions. In such a case,

$\frac{\partial}{\partial k}\omega_j(x) = 0$. $\frac{\partial}{\partial k} f_Y(x|X, \pi(k), t, 1)$ exists for appropriately chosen information structures. For example, one can verify that if $f_Y(x|X, \pi(k)) \sim Beta(yk, zk)$, then $\frac{\partial}{\partial k} f_Y(x|X, \pi(k), t, 1)$ exists and is bounded, which also takes care of the second part of condition 1.

Regarding condition 2, from assumption A3 it holds whenever $|\pi_0 - \pi^*| \leq \overline{\pi}$. Otherwise, we can have $\frac{\partial}{\partial k}\mu_{j,1,1} \geq 0$ on one side. This complicates the proof but should not change the conclusion. Suppose this is the case for buyers so, $\uparrow k \Rightarrow \uparrow \mu_{B,1,1} \Rightarrow \downarrow (1 - V_{B,1,1}) = r_{B,1,1}$. In such cases, the effect on $\mu_{S,1,1}$ is increased for the same reason so that $\uparrow k \Rightarrow \downarrow\downarrow \mu_{S,1,1} \Rightarrow \downarrow\downarrow V_{S,1,1} = r_{S,1,1}$ so it is reasonable to expect that price dispersion is still increasing with k.

5.1.5 Prior Precision and Price Dispersion

Several existing models allow one to study the effects of the precision of the prior on transaction price dispersion. These include Burdett and Vishwanath (1988), Anenberg (2016) and Glaeser and Nathanson (2017). In all three, a more precise prior reduces the dispersion in beliefs caused by private information. Lower dispersion in beliefs translates into lower dispersion in transaction prices. The exact details differ slightly, but all three share the feature that only one side of the market faces meaningful uncertainty and makes strategic decisions.[17]

Proposition 2.1 If

1. $T_B(k_1) = T_B(k_0)$.
2. Only buyers make offers.
3. An offer must equal an agent's true reservation price $o_{B,t,1} = r_{B,t,1}$.
4. Sellers always accept.

Then the dispersion of transaction prices is lower under k_1 compared with k_0.

That proposition 2.1 holds is immediate from assumption A1. The above proposition also holds if the roles of buyers and sellers are reversed. In the present setting, higher k (more precise prior) reduces the dispersion in beliefs, just as in the works cited above. This effect is unambiguous on the intensive margin by assumption A1. The range of beliefs can increase if a higher k also causes

[17] This brings to mind the famous Rothschild (1974) critique of early search models that did not model the side making the bids. I should note that none of those three papers analyzes the effects of the prior precision on price dispersion. I was able to verify that my claim about Glaeser and Nathanson (2017) holds only when changed their model slightly by assuming demand follows a random walk and that time is discrete.

postponed quitting. However, because the threshold k's for postponed quitting (k^*) are discrete, for any given k_0, one can always find k_1 such that $k_0 < k_1 < k^*$. Therefore, when the present setting is reduced to a one sided version by letting only buyers offer and forcing their bids to either equal their true valuations (as in Glaeser and Nathanson, 2017), price dispersion in transaction prices decreases with public information (if public information does not change quitting decisions).

Proposition 2 shows that different conclusions may be reached in a two sided uncertainty, strategic bargaining setting. FTE shuts down the transmission of within side belief dispersion to transaction price dispersion which highlights the impact of public information on dispersion between buyers and sellers on the bargaining margin.

5.1.6 Market Makers

Market makers are market participants that act both as buyers and as sellers at the same time and on a large scale. They have long been a feature of financial markets. A recent study (Buchak et al, 2020) views iBuyers as playing a similar role in the housing market. In this section, I introduce them into the current model to help illustrate the relation between the above price dispersion result (Proposition 2) and existing results concerning the bid-ask spread. In particular, I show that the bid-ask spread decreases with public information, corroborative the result of Glosten and Milgrom (1985) in a new setting. I also show that when the probability of meeting a market maker is not too large, total price variance may still increase with public information.

Duffie et al (2005) present a model of a market where participants (buyers and sellers) randomly meet with each other and with market makers. The market makers are assumed to have no holding costs or inventory risk so they are concerned only with maximizing profits from each meeting. This conception of market markers is particularly relevant to decentralized markets.

Now suppose that market makers as in Duffie et al (2005) are introduced into the present model so that with some known probability $\rho > 0$, a participant encounters a market maker. I am assuming this event is realized after entry and quitting but before meetings between buyers and sellers (between stage 2 and 3 in the order of events). The market maker faces a similar trade-off to other participants about whether to make an offer everyone accepts or a riskier offer. Because it has unlimited inventory, its only concern is to choose the bid that maximizes immediate profits. Therefore, if k is not too small so that heterogeneity is not too large, it is optimal for the market market to make the offer everyone would accept. I assume k is small enough both for the above and for FTE to be the unique perfect

Bayesian equilibrium.[18]

Proposition 2.2

1. If market makers are introduced, they offer to trade $r_{j,0,0}$ when facing a an agent of type $(j, t, 0)$.

2. If in addition $|\pi_0 - \pi^*| \leq \overline{\pi}$, then the bid-ask spread $r_{B,0,0} - r_{S,0,0}$ is smaller for k_1 compared with k_0 for all $k_0 < k_1 < \infty$.

3. If in addition, ρ is not too large and the conditions of Proposition 2 hold, then the variance of transaction prices is increasing in k.

Proof: Consider a meeting between a market marker and a seller. Because the probability of meeting a market marker is known, this event does not provide information about the state of the world to the seller. Because we are in FTE, it must be the case that the seller has never met a buyer. Therefore, the most optimistic seller has a history $(0, 0)$. The market marker maximizes its profit by offering $r_{S,0,0}$ (this is conditional on k not too small as explained above) and the seller accepts regardless of her true history (because all sellers in the market at this stage have reservations weakly below $r_{S,0,0}$). A meeting between a market marker and a buyer is analogous.

Therefore, with market makers, the equilibrium price distribution has 4 mass points: two from meetings including market makers and two from meetings between buyers and sellers. These are: $r_{S,0,0}, r_{S,1,1}, r_{B,1,1}$ and $r_{B,0,0}$. Note that $r_{S,1,1} > r_{S,0,0}$ because $r_{S,1,1}$ is based on 1 positive and 0 negative signals. $r_{B,1,1} < r_{B,0,0}$ for the same reason. Therefore, we get $r_{S,0,0} < r_{S,1,1} < r_{B,1,1} < r_{B,0,0}$.

Now, compare two markets that are identical except for a difference in k with $k_0 < k_1 < \infty$. From assumption A3, if $|\pi_0 - \pi^*| \leq \overline{\pi}$ then $\mu_{j,t,0}(k_1) > \mu_{j,t,0}(k_0)$, which implies $d_{j,0,0}(k_1) > d_{j,0,0}(k_0)$ and $g_{j,0,0}(k_1) < g_{j,0,0}(k_0)$ so that $V_{j,0,0}(k_1) > V_{j,0,0}(k_0)$. The latter in turn implies $r_{S,0,0}(k_1) > r_{S,0,0}(k_0)$ and $r_{B,0,0}(k_1) < r_{B,0,0}(k_0)$ (recall that $r_{B,t,m} = 1 - V_{B,t,m}$). Therefore, $r_{B,0,0}(k_1) - r_{S,0,0}(k_1) < r_{B,0,0}(k_0) - r_{S,0,0}(k_0)$ so that the price range is decreasing in k. At the same time, it is still the case that for ρ small enough, the proof of $\frac{\partial}{\partial k}(r_{B,1,1} - r_{S,1,1}) > 0$ is essentially unchanged. Small enough ρ also guarantees that the prices $r_{S,0,0}, r_{B,0,0}$ are rare relative to $r_{B,1,1}, r_{S,1,1}$ so that the effect of $\frac{\partial}{\partial k}(r_{B,1,1} - r_{S,1,1}) > 0$ dominates. In this case, higher k causes higher price variance but a lower price range.

[18] For the purpose of the arguments that follow, ρ can be thought of as being arbitrarily close to 0 so that all the arguments for existence and uniqueness of FTE are essentially unchanged. However, there is not reason to suppose a larger ρ would invalidate those results, only that it might make them more difficult to derive.

QED

5.2 Welfare

Improved public information may affect welfare through its effect on quitting decisions. Postponed quitting increases the total population. Increased population must increase total search costs which reduces total welfare, but may also increase total matches which increases total welfare. The sign of the total effect on welfare depends on the relative size of these two effects. I focus here on the cases where participation is objectively beneficial for both buyers and sellers ($\omega_j u_j > c_j, \forall j$) and where both sides in fact participate under the two levels of k that I examine. I discuss other cases later on.

The net surplus generated by this market in each period is given by

$$W = M(N_S, N_B) - N_S \cdot c_S - N_B \cdot c_B \tag{5}$$

Where $N_j = \sum_{t=0}^{T_j} N_{j,t}$ and $N_{j,t}$ is the mass of agents of type (j,t) after entry and quitting decisions but before meetings are determined.[19] $M(,)$ is the aggregate matching function which determines the mass of meetings given masses of agents.

Equation (5) captures the idea that each meeting in FTE generates 1 unit of gross surplus and that total net surplus equals the number of meetings (technically, the **mass** of meetings) minus total search costs paid by all participants. Therefore, if $k_0 < k_1$, the effect of public information on welfare is given by:

$$\Delta W \equiv W(k_1) - W(k_0) = \Delta M + \sum_{j \in \{S,B\}} \Delta N_j \cdot c_j \tag{6}$$

Where $\Delta M \equiv M(N_S(k_1), N_B(k_1)) - M(N_S(k_0), N_B(k_0))$ and $\Delta N_j \equiv N_j(k_1) - N_j(k_0)$.

First, note that in FTE k affects aggregate matches M only through the populations N_j. Second, the term ΔN_j can be decomposed into two effects: 1) is the primary (marginal) effect which changes population by eliminating or adding types of agents who participate in the market by changing the quitting period T (denoted by $\Delta_1 N_j$) and 2) the secondary (infra-marginal) effect which changes population by changing the probability of reaching a given tenure which is caused by changing the per

[19] I am omitting the subscript m without loss of generality because quitting decisions are made before matching is announced so that in FTE, at this stage, $m = 0$ for all agents.

period matching rate (denoted by $\Delta_2 N_j$). So that $\Delta N_j = \Delta_1 N_j + \Delta_2 N_j$. This implies that $\Delta_1 N_j$ is the sole primary effect in ΔW because both the effect on ΔM and on $\Delta_2 N_j$ operate are secondary and operate only through $\Delta_1 N_j$.

Define the model set, $\mathcal{M}_i$ which includes a particular choice for state space X, realized public signal π_0, matching function $M(N_S, N_B)$, buyer entry rate function $N_{B,0}(x)$ and parameter vector $(N_{S,0}, c_{B,} c_S, \delta_B, \delta_S, k_i)$. For each $\mathcal{M}_i$ one can directly calculate the full trade equilibrium value function $V_{j,t,m}$ and by doing so find $T_j(k)$ as described in 4.5. The mapping from $\mathcal{M}_i$ to $T_j(k_i)$ is not particularly informative so I work directly with $T_j(k_i)$ by studying separately the different directions in which $T_j(k)$ can be affected by k.

In what follows, I focus on the cases where the effect of k on T_j is non-negative, which is always the case when $|\pi_0 - \pi^*| \leq \overline{\pi}$. This clarifies the analysis and the resulting insights can also be applied to the other cases without much modification. This implies that $\Delta T_j = T_j(k_1) - T_j(k_0) \geq 0$ and so $\Delta_1 N_j = \sum_{t=0}^{T_j(k_1)} N_{j,t} - \sum_{t=0}^{T_j(k_0)} N_{j,t} \geq 0$. Therefore, we need to consider consider four cases:

1. $\Delta T_S = \Delta T_B = 0$
2. $\Delta T_S = \Delta T_B > 0$
3. $\Delta T_B > \Delta T_S = 0$
4. $\Delta T_S > \Delta T_B = 0$

In what follows I consider the cases where if the market is unbalanced so that $N_{S,0} \neq N_{B,0}(x)$ then $N_{S,0} > N_{B,0}(x)$ so it is a buyer's market. As will be made clear, in terms of the welfare effects of k, buyers and sellers are interchangeable, but it matters whether T increases on the short or on the long side of the market.

Proposition 3

1. If $\Delta T_S = \Delta T_B = 0$, then $\Delta W = 0$
2. If ($\Delta T_S = \Delta T_B > 0$ or $\Delta T_B > \Delta T_S = 0$) and $\omega_j(k_1) \geq \omega_j(k_0)$ $\forall j$, then $\Delta W > 0$
3. $\Delta T_S > \Delta T_B = 0$ and $\Delta M \leq 0$, then $\Delta W < 0$

The proof proceeds on a case by case basis

5.2.1 Case 1

If $\Delta T_S = \Delta T_B = 0$ then $\Delta N_j = 0$ which implies that $\Delta W = 0$. This occurs when on both sides $V^c_{j,T_j,0}(k_0) \leq V^c_{j,T_j,0}(k_1) < 0$. The first inequality is always true because T_j negative signals have a weaker effect on the posterior for higher k. The second inequality can occur when the increase in k from k_0 to k_1 is insufficient to tip $V^c_{j,T_j,0}(k)$ over the zero threshold. Whether this is or is not the case depends on the choice of $\mathcal{M}_0, \mathcal{M}_1$. For a given k_0, there always exists k_1 sufficiently large so that $T_j(k_1) > T_j(k_0)$ for all j. In what follows, I focus on such choices of k_0, k_1.

5.2.2 Cases 2 and 3

Suppose $\Delta T_S = \Delta T_B > 0$ and matching chances are unchanged so that $\Delta_1 N_j > 0$ for both $j = \{B, S\}$. In this case, $\Delta W > 0$ whenever $\omega_j u_j > c_j$ for both $j = \{B, S\}$. To see why, note that in this case there is no externality and the only agents affected are those that postpone their exit. From an objective perspective, reflected by ω_j in the above term, participation in a single round is beneficial whenever the expected benefit exceeds the search cost.

The above also implies that if matching chances increase on both sides, $\Delta W > 0$ as well. If matching chances increase on one side but decrease on the other, $Sign\{\Delta W\}$ will depend on the relative sizes of the parameters including the elasticity of the matching function w.r.t population.

As an example of a case where quitting is postponed and matching chances do not decrease on either side, one may consider a scenario where the effect of postponed quitting on population is symmetric ($\Delta N_S = \Delta N_B$) and the matching function has non-decreasing returns to scale (non-DRS). $\Delta N_S = \Delta N_B$ occurs for example when $N_{S,0} = N_{B,0}$, $c_S = c_B, \delta_S = \delta_B$. Note that these conditions are not necessary for the above but only sufficient. Also, if matching is DRS, then $Sign\{\Delta W\}$ depends on the relative strength of DRS and the objective benefit from search.

If $\Delta T_B > \Delta T_S = 0$ this implies that the primary effect is to increase population on the short side but not on the long side. This tends to increase long side matching rate more than it decreases short side matching chances. An especially clear cut example is the Leontief matching function $M_L(N_S, N_B) \equiv A \cdot Min\{N_S, N_B\}$, where short side matching chances do not decrease at all and the number of matches increases proportionately with short side population.

5.2.3 Case 4

Suppose $\Delta T_S > \Delta T_B = 0$ so that population increases only on the long side. This is the reverse of case 3, but here the Leontief matching function implies that total matches do not increase at all, so the only effect of increased population is increased congestion. More generally, if no extra matches are generated by the increased population, welfare must decline.

5.2.4 Remarks

1. k may also affect entry decisions. Entering agents are all ex-ante identical except for their role (buyer or seller), so the initial entry decision of paying the first search cost is the same for everyone on each side. From an objective standpoint, entry is beneficial whenever $\omega_j(x)u_j - c_j \geq 0$. Since agents do not observe x, they choose to pay the first search cost whenever $\mu_{j,0,0}(\pi_0, k)u_j - c_j \geq 0$. A market exists only if $\mu_{S,0,0}(\pi_0, k)u_S - c_S \geq 0$ and $\mu_{B,0,0}(\pi_0, k)u_B - c_B \geq 0$. Increasing k can reduce $\mu_{S,0,0}$ or $\mu_{B,0,0}$ which can eliminate the market. This can occur regardless $Sign\{\omega_j(x)u_j - c_j\}$ on both sides. Note that because π is an unbiased signal of x, $E_\pi[\mu_{j,0,0}(\pi_0, k)] = \omega_j(x)$ so a larger k does not increase or decrease the accuracy of beliefs (prior to learning) or the probability of participation.

2. Proposition 3 is reflected in observed search durations. To simplify the discussion, I am again assuming that participation is objectively beneficial and the market exists under both k_0 and k_1. First, note that maximal search durations are higher whenever quitting occurs later (by definition). Second, whenever later quitting results in a larger population on side j, this implies longer average time on the market on side j. This is because entry is unaffected by k so a higher population must result from reduced exit and thus longer average time on the market. Therefore, if k results in longer maximal search durations and **no** longer average search duration, this implies that matching rates increased, which implies welfare has increased. The above implies that statistics of search durations can be used to infer the welfare effects of public information by a partially informed social planner. In particular, if a higher k is not associated with higher average search durations on either side, its effect on welfare is non-negative. If, in addition, higher k results in higher maximal search durations, it's effect must be welfare improving. This insight is especially important in the housing market where search durations can often be measured with reasonable accuracy only in the events where the search yielded a successful transaction.

6 Conclusion

I study how public information and sequential learning interact in the context of a bilateral market. A richer state space compared with similar previous studies allows me to perform comparative statics with respect to the precision of public information. I provide general conditions for a declining belief path during search and study a large class of economies where public information mitigates this decline. This has implications for premature quitting, declining reservation prices and search durations. More importantly, I characterize a novel mechanism through which public information improves or diminishes welfare. Furthermore, allowing for two sided uncertainty and hidden continuation values provides new insights about the effects of public information on price dispersion.

Further research should proceed along two avenues: First, sequential learning and public information about uncertain market conditions should be embedded in other market structures than the one treated here. Second, the predictions of this model for price dispersion, search durations and declining reservation prices should be tested empirically.

References

Afrouzi, Hassan, Spencer Yongwook Kwon, Augustin Landier, Yueran Ma, and David Thesmar. Overreaction and working memory. No. w27947. National Bureau of Economic Research, 2020.

Amador, Manuel, and Pierre-Olivier Weill. "Learning from prices: Public communication and welfare." Journal of Political Economy 118, no. 5 (2010): 866-907.

Anenberg, Elliot. "Information frictions and housing market dynamics." International Economic Review 57, no. 4 (2016): 1449-1479.

Anenberg, Elliot, and Steven Laufer. "A more timely house price index." Review of Economics and Statistics 99, no. 4 (2017): 722-734.

Angeletos, George-Marios, and Alessandro Pavan. "Efficient use of information and social value of information." Econometrica 75, no. 4 (2007): 1103-1142.

Baye, Michael R., John Morgan, and Patrick Scholten. "Information, search, and price dispersion." Handbook on economics and information systems 1 (2006): 323-375.

Ben-Shahar, D., & Golan, R. (2019). Improved information shock and price dispersion: A natural experiment in the housing market. Journal of Urban Economics, 112, 70-84.

Blouin, Max R., and Roberto Serrano. "A decentralized market with common values uncertainty: Non-steady states." The Review of Economic Studies 68, no. 2 (2001): 323-346.

Blundell, Richard, John Ham, and Costas Meghir. "Unemployment, discouraged workers and female labour supply." Research in Economics 52, no. 2 (1998): 103-131.

Bordalo, Pedro, Nicola Gennaioli, Yueran Ma, and Andrei Shleifer. "Overreaction in macroeconomic expectations." American Economic Review 110, no. 9 (2020): 2748-82.

Brown, Meta, Christopher J. Flinn, and Andrew Schotter. "Real-time search in the laboratory and the market." American Economic Review 101, no. 2 (2011): 948-74.

Buchak, Greg, Gregor Matvos, Tomasz Piskorski, and Amit Seru. "iBuyers: Liquidity in Real Estate Markets?." Available at SSRN 3616555 (2020).

Burdett, K., & Vishwanath, T. (1988). Declining reservation wages and learning. The Review of Economic Studies, 55(4), 655-665.

da Silveira, Rava Azeredo, Yeji Sung, and Michael Woodford. Optimally imprecise memory and biased forecasts. No. w28075. National Bureau of Economic Research, 2020.

Diamond, Peter A. "A model of price adjustment." Journal of economic theory 3, no. 2 (1971): 156-168.

Duffie, Darrell, Nicolae Gârleanu, and Lasse Heje Pedersen. "Over-the-counter markets." Econometrica 73, no. 6 (2005): 1815-1847.

Duffie, Darrell, Semyon Malamud, and Gustavo Manso. "Information percolation with equilibrium search dynamics." Econometrica 77, no. 5 (2009): 1513-1574.

Duffie, Darrell, Semyon Malamud, and Gustavo Manso. "Information percolation in segmented markets." Journal of Economic Theory 153 (2014): 1-32.

Duffie, Darrell, Piotr Dworczak, and Haoxiang Zhu. "Benchmarks in search markets." The Journal of Finance 72, no. 5 (2017): 1983-2044.

Dye, Ronald A. "Disclosure of nonproprietary information." Journal of accounting research (1985): 123-145.

Eckwert, Bernhard, and Itzhak Zilcha. "Incomplete risk sharing arrangements and the value of information." Economic Theory 21, no. 1 (2003): 43-58.

Fajgelbaum, Pablo D., Edouard Schaal, and Mathieu Taschereau-Dumouchel. "Uncertainty traps." The Quarterly Journal of Economics 132, no. 4 (2017): 1641-1692.

Frenkel, Sivan, Ilan Guttman, and Ilan Kremer. "The effect of exogenous information on voluntary disclosure and market quality." Journal of Financial Economics (2020).

Gale, Douglas. "Bargaining and competition part I: characterization." Econometrica: Journal of the Econometric Society (1986): 785-806.

Gale, Douglas. Strategic foundations of general equilibrium: dynamic matching and bargaining games. Cambridge University Press, 2000.

Genesove, David, and Christopher J. Mayer. "Equity and Time to Sale in the Real Estate Market." The American Economic Review (1997): 255-269.

Genesove, David, and Lu Han. "Search and matching in the housing market." Journal of Urban Economics 72, no. 1 (2012): 31-45.

Glaeser, Edward L., and Charles G. Nathanson. "An extrapolative model of house price dynamics." Journal of Financial Economics 126, no. 1 (2017): 147-170.

Glosten, Lawrence R., and Paul R. Milgrom. "Bid, ask and transaction prices in a specialist market with heterogeneously informed traders." Journal of financial economics 14, no. 1 (1985): 71-100.

Golosov, Mikhail, Guido Lorenzoni, and Aleh Tsyvinski. "Decentralized trading with private information." Econometrica 82, no. 3 (2014): 1055-1091.

Gronau, Reuben. "Information and frictional unemployment." The American Economic Review

(1971): 290-301.

Hirshleifer, Jack "The private and social value of information and the reward to incentive activity". American Economic Review (1971): 561–574

Hirshleifer, David. "Presidential Address: Social Transmission Bias in Economics and Finance." The Journal of Finance (2020).

Jensen, Robert. "The digital provide: Information (technology), market performance, and welfare in the South Indian fisheries sector." The quarterly journal of economics 122, no. 3 (2007): 879-924.

Jovanovic, Boyan, and Sai Ma. "Uncertainty and growth disasters." Review of Economic Dynamics (2021).

Jovanovic, Boyan, and Albert J. Menkveld. "Equilibrium Bid-Price Dispersion." Available at SSRN 2463066 (2020).

Kagel, John H., and Dan Levin. "The winner's curse and public information in common value auctions." The American economic review (1986): 894-920.

Kagel, John H., and Alvin E. Roth, eds. The handbook of experimental economics, volume 2: the handbook of experimental economics. Princeton university press, 2016.

Kotova, Nadia, and Anthony Lee Zhang. "Search Frictions and Idiosyncratic Price Dispersion in the US Housing Market." Available at SSRN 3386353 (2020).

Kotsenko, Nikita. "Information, Distribution of Prices and Time on the Market". Alrov Discussion Paper, available at https://alrov.tau.ac.il/?page_id=2182 (2018)

Kozlowski, Julian, Laura Veldkamp, and Venky Venkateswaran. "The tail that keeps the riskless rate low." NBER Macroeconomics Annual 33, no. 1 (2019): 253-283.

Lauermann, Stephan, Wolfram Merzyn, and Gábor Virág. "Learning and price discovery in a search market." The Review of Economic Studies 85, no. 2 (2017): 1159-1192.

Levin, Dan, James Peck, and Asen Ivanov. "Separating Bayesian updating from non-probabilistic reasoning: An experimental investigation." American Economic Journal: Microeconomics 8, no. 2 (2016): 39-60.

Levitt, Steven D., and Chad Syverson. "Market distortions when agents are better informed: The value of information in real estate transactions." The Review of Economics and Statistics 90, no. 4 (2008): 599-611.

Ljungqvist, Lars, and Thomas J. Sargent. "Understanding European unemployment with matching and search-island models." Journal of Monetary Economics 54, no. 8 (2007): 2139-2179.

Majumdar, Dipjyoti, Artyom Shneyerov, and Huan Xie. "An optimistic search equilibrium." Review of Economic Design 20, no. 2 (2016): 89-114.

Merlo, Antonio, and Francois Ortalo-Magne. "Bargaining over residential real estate: evidence from England." Journal of urban economics 56, no. 2 (2004): 192-216.

Merlo, Antonio, François Ortalo-Magné, and John Rust. "The home selling problem: Theory and evidence." International Economic Review 56, no. 2 (2015): 457-484.

Milgrom, Paul R., and Robert J. Weber. "A theory of auctions and competitive bidding." Econometrica: Journal of the Econometric Society (1982): 1089-1122.

Morris, Stephen, and Hyun Song Shin. "Social value of public information." American economic review 92, no. 5 (2002): 1521-1534.

Ottaviani, Marco, and Andrea Prat. "The value of public information in monopoly." Econometrica 69, no. 6 (2001): 1673-1683.

Pavan, Alessandro, and Xavier Vives. "Information, coordination, and market frictions: An introduction." Journal of Economic Theory 158, no. Part B (2015): 407-426.

Pennerstorfer, Dieter, Philipp Schmidt-Dengler, Nicolas Schutz, Christoph Weiss, and Biliana Yontcheva. "Information and price dispersion: Theory and evidence." International Economic Review 61, no. 2 (2020): 871-899.

Petrongolo, Barbara, and Christopher A. Pissarides. "Looking into the black box: A survey of the matching function." Journal of Economic literature 39, no. 2 (2001): 390-431.

Piazzesi, Monika, Martin Schneider, and Johannes Stroebel. "Segmented housing search." American Economic Review 110, no. 3 (2020): 720-59.

Rogerson, Richard, and Robert Shimer. "Search in macroeconomic models of the labor market." In Handbook of Labor Economics, vol. 4, pp. 619-700. Elsevier, 2011.

Rothschild, Michael. "Searching for the lowest price when the distribution of prices Is unknown: a summary." In Annals of Economic and Social Measurement, Volume 3, number 1, pp. 293-294. NBER, 1974.

Rubinstein, Ariel, and Asher Wolinsky. "Equilibrium in a market with sequential bargaining." Econometrica: Journal of the Econometric Society (1985): 1133-1150.

Shneyerov, Artyom, and Adam CL Wong. "Price discovery in a matching and bargaining market with aggregate uncertainty." Games and Economic Behavior 124 (2020): 183-206.

Smith, Vernon L., Gerry L. Suchanek, and Arlington W. Williams. "Bubbles, crashes, and endogenous expectations in experimental spot asset markets." Econometrica: Journal of the Econometric Society (1988): 1119-1151.

Stigler, George J. "The economics of information." Journal of political economy 69, no. 3 (1961): 213-225.

Stoll, Hans R. "Presidential address: friction." The Journal of Finance 55, no. 4 (2000): 1479-1514.

Van Nieuwerburgh, Stijn, and Laura Veldkamp. "Learning asymmetries in real business cycles." Journal of monetary Economics 53, no. 4 (2006): 753-772.

Varian, Hal R. "A model of sales." The American economic review 70, no. 4 (1980): 651-659.

Wolinsky, Asher. "Information revelation in a market with pairwise meetings." Econometrica: Journal of the Econometric Society (1990): 1-23.

Zeira, Joseph. "Informational cycles." The Review of Economic Studies 61, no. 1 (1994): 31-44.

Zeira, Joseph. "Informational overshooting, booms, and crashes." Journal of Monetary Economics 43, no. 1 (1999): 237-257.

A Continuation Value

If bargaining breaks down for an agent of type $(j, t, 1)$, she has a probability δ_j of not exiting at the end of the period. Assuming $t \leq T_j - 1$ so that the agent will choose to continue, she will pay c_j, and then have a probability of $\mu_{j,t,1}$ to exit with utility u_j and a probability $(1 - \mu_{j,t,1})$ to continue to the next period where her expected utility will be $V_{j,t+1,1}$:

$$V_{j,t,1} = \delta_j \mu_{j,t,1} u_j + \delta_j (1 - \mu_{j,t,1}) \cdot V_{j,t+1,1} - \delta_j \cdot c_j$$

This is true for all t, in particular, for $t + 1$:

$$V_{j,t+1,1} = \delta_j \mu_{j,t+1,1} u_j + \delta_j (1 - \mu_{j,t+1,1}) \cdot V_{j,t+2,1} - \delta_j \cdot c_j$$

Therefore, by repeatedly substituting $V_{j,t+1,1}, V_{j,t+2,1}, ...$ one can obtain:

$$\begin{aligned} V_{j,t,1} = {} & \delta_j \mu_{j,t,1} u_j + \delta_j (1 - \mu_{j,t,1}) \cdot \delta_j \mu_{j,t+1,1} u_j \\ & + \delta_j (1 - \mu_{j,t,1}) \cdot \delta_j (1 - \mu_{j,t+1,1}) \cdot \delta_j \mu_{j,t+2,1} u_j \\ & - \delta_j \cdot c_j - \delta_j (1 - \mu_{j,t,1}) \cdot \delta_j \cdot c_j - \delta_j (1 - \mu_{j,t,1}) \cdot \delta_j (1 - \mu_{j,t+1,1}) \cdot \delta_j \cdot c_j \\ & + \delta_j (1 - \mu_{j,t,1}) \cdot \delta_j (1 - \mu_{j,t+1,1}) \cdot \delta_j (1 - \mu_{j,t+2,1}) \cdot V_{j,t+3,1} \end{aligned}$$

$$
\begin{aligned}
V_{j,t,1} &= \delta_j(\mu_{j,t,1}u_j - c_j) + \delta_j^2(1-\mu_{j,t,1})(\mu_{j,t+1,1}u_j - c_j) \\
&+ \delta_j^3(1-\mu_{j,t,1})(1-\mu_{j,t+1,1})(\mu_{j,t+2,1}u_j - c_j) \\
&+ \delta_j^3(1-\mu_{j,t,1})(1-\mu_{j,t+1,1})(1-\mu_{j,t+2,1}) \cdot V_{j,t+3,1} \\
&= \delta_j(\mu_{j,t,1}u_j - c_j) + \sum_{s=2}^{3} \delta_j^s \prod_{i=t}^{t+s-2} (1-\mu_{j,i,1})(\mu_{j,t+s-1,1}u_j - c_j) + \delta_j^3 V_{j,t+3,1} \\
&= \delta_j(\mu_{j,t,1}u_j - c_j) + \sum_{s=2}^{T_j - t} \delta_j^s \prod_{i=t}^{t+s-2} (1-\mu_{j,i,1})(\mu_{j,t+s-1,1}u_j - c_j) \\
&= [\delta_j \mu_{j,t,1} + \sum_{s=2}^{T_j - t} \delta_j^s \prod_{i=t}^{t+s-2} (1-\mu_{j,i,1})\mu_{j,t+s-1,1}]u_j \\
&- [\delta_j + \sum_{s=2}^{T_j - t} \delta_j^s \prod_{i=t}^{t+s-2} (1-\mu_{j,i,1})]c_j \\
&\equiv d_{j,t,1}(T_j) \cdot u_j(T_j, T_{-j}) - g_{j,t,1}(T_j) \cdot c_j
\end{aligned}
$$

Note that $\delta^{T_j} V_{j,t+(T_j-t+1),1}$ is dropped after T_j is introduced because the agent exits at the start of period T_j, so from then on her continuation utility is zero.

B Prices

Now note that in a full trade equilibrium:

$$u_S = \frac{1}{2} r_{B,1} + \frac{1}{2} r_{S,1}$$

$$u_B = \frac{1}{2}(1 - r_{B,1}) + \frac{1}{2}(1 - r_{S,1})$$

Also,

$$r_{S,1} = V_{S,1,1}$$

$$r_{B,1} = 1 - V_{B,1,1}$$

Therefore, we can write:[20]

$$V_{S,1,1} = d_{S,1,1} \cdot \frac{1}{2}[r_{B,1} + r_{S,1}] = d_{S,1,1} \cdot \frac{1}{2}[1 - V_{B,1,1} + V_{S,1,1}]$$

$$V_{B,1,1} = d_{B,1} \cdot \frac{1}{2}[2 - r_{B,1} - r_{S,1}] = d_{B,1,1} \cdot \frac{1}{2}[1 + V_{B,1,1} - V_{S,1,1}]$$

So we have 2 equations and 2 unknowns $(V_{S,1,1}, V_{B,1,1})$ so we can solve for them in terms of $d_{S,1,1}, d_{B,1,1}$

$$V_{S,1} = \frac{d_{S,1,1}(1 - d_{B,1,1})}{2 - d_{B,1,1} - d_{S,1,1}}$$

$$V_{B,1} = \frac{d_{B,1,1}(1 - d_{S,1,1})}{2 - d_{B,1,1} - d_{S,1,1}}$$

Therefore,

$$r_{S,1} = \frac{d_{S,1,1}(1 - d_{B,1,1})}{2 - d_{B,1,1} - d_{S,1,1}}$$

$$r_{B,1} = 1 - \frac{d_{B,1,1}(1 - d_{S,1,1})}{2 - d_{B,1,1} - d_{S,1,1}}$$

And note that

$$r_{B,1} > r_{S,1} \Leftrightarrow 1 - \frac{d_{B,1,1}(1 - d_{S,1,1})}{2 - d_{B,1,1} - d_{S,1,1}} > \frac{d_{S,1,1}(1 - d_{B,1,1})}{2 - d_{B,1,1} - d_{S,1,1}}$$
$$\Leftrightarrow 1 - d_{B,1,1} - d_{S,1,1} + d_{B,1,1}d_{S,1,1} > 0$$

is always true because $d_{j,1,1} \in (0, 1)$

This also gives us

$$u_S = \frac{1}{2}(\frac{d_{B,1,1} + d_{S,1,1} - 2d_{S,1,1}d_{B,1,1}}{2 - d_{B,1,1} - d_{S,1,1}})$$

[20] I am assuming $c_j = 0$, the derivation in case $c_j > 0$ is analogous.

$$u_B = \frac{1}{2}[2 - \frac{d_{B,1,1} + d_{S,1,1} - 2d_{S,1,1}d_{B,1,1}}{2 - d_{B,1,1} - d_{S,1,1}}]$$

And that any continuation value $V_{B,t,m}$ can be written as only a function of $d_{B,t,m}, d_{S,t,m}$:

$$\begin{aligned} V_{B,t,m} &= d_{B,t,m} \cdot u_B \\ &= d_{B,t,m}\frac{1}{2}(2 - \frac{d_{B,1,1} + d_{S,1,1} - 2d_{S,1,1}d_{B,1,1}}{2 - d_{B,1,1} - d_{S,1,1}}) \end{aligned}$$

$$V_{S,t,m} = d_{S,t,m}\frac{1}{2}(\frac{d_{B,1,1} + d_{S,1,1} - 2d_{S,1,1}d_{B,1,1}}{2 - d_{B,1,1} - d_{S,1,1}})$$

A similar result can be obtained in the same way for the case with $c_j > 0$.

C Steady State

$N_{j,t}(x;l)$ is the number of agents with role $j = \{B, S\}$ and tenure $t = \{0, 1, 2, ...\}$ in period l under state x.

I am assuming that the entry rates $N_{j,0}(x)$, are large enough so that the residual population in all tenures $N_{j,t}(x;l)$ is large enough so that the law of large numbers applies. In particular: if each agent in $N_{j,t}(x;l)$ experiences some event with probability δ in each period, then in each period the number of agents in $N_{j,t}(x;l)$ that experience this event equals $\delta \cdot N_{j,t}(x;l)$. If she finds it more convenient, the reader may suppose that $N_{j,1}(x)$ is the mass of a continuum of agents.

Claim: For any initial stocks of agents $N_{j,t}(x;l) > 0$: $lim_{n\to\infty} N_{j,t}(x;l+n) = N_{j,t}(x)$

Proof:

In period 1, we have entry of $N_{S,0}(x)$, $N_{B,0}(x)$, and this is the entire population. Then, $M(N_{S,0}(x), N_{B,0}(x)) = p_1$ pairs are formed. Then out of the remaining buyers and sellers, a fraction $1 - \delta$ exits exogenously so that tenure 1 seller population at the start of period 2 is: $N_{j,1}(x;2) = \delta(N_{j,0}(x) - p_1)$.

In period 2, $M(N_S(x;2), N_B(x;2)) = p_2$, where $N_j(x;l) \equiv \sum_{t=1}^{\infty} N_{j,t}(x;l)$[21]. Part of them come from tenure 0 and part from tenure 1. So tenure 1 population in period 3 is: $N_{j,1}(x;3) = \delta(N_{j,0}(x) -$

[21] I am assuming here that $c_j = 0$ so that there is no quitting. I relax this assumption later on.

$p_2 \frac{N_{j,0}(x)}{N_j(x;2)}$) because tenure 1 is populated by everyone from tenure 0 who hasn't exited exogenously or matched and $p_2 \frac{N_{j,0}(x)}{N_j(x;2)}$ is the number of meetings in period 2 that were drawn from tenure 0.

Similarly: $N_{j,2}(x;3) = \delta(N_{j,1}(x;2) - p_2 \frac{N_{j,1}(x;2)}{N_j(x;2)})$.

Suppose that we are in a period l large enough so that all tenures are populated. Then

$$\begin{aligned} N_{j,t}(x;l+1) &= \delta(N_{j,t-1}(x;l) - p_l \frac{N_{j,t-1}(x;l)}{N_j(x;l)}) \\ &= \delta N_{j,t-1}(x;l)(1 - p_l \frac{1}{N_j(x;l)}) \end{aligned}$$

So that

$$\begin{aligned} N_j(x;l) &= \sum_{t=0}^{\infty} N_{j,t}(x;l) \\ &= \sum_{t=0}^{\infty} \delta N_{j,t-1}(x;l-1)(1 - p_{l-1} \frac{1}{N_j(x;l-1)}) \\ &= \sum_{t=0}^{\infty} \delta^2 N_{j,t-2}(x;l-2)(1 - p_{l-1} \frac{1}{N_j(x;l-1)})(1 - p_{l-2} \frac{1}{N_j(x;l-2)}) \\ &= \sum_{t=0}^{\infty} \delta^s N_{j,t-s}(x;l-s) \prod_{r=0}^{s} (1 - p_{l-r} \frac{1}{N_j(x;l-r)}) \\ &= N_{j,0}(x) \sum_{t=0}^{\infty} \delta^{t-1} \prod_{r=0}^{(t-1)} (1 - p_{l-r} \frac{1}{N_j(x;l-r)}) \end{aligned}$$

The next to last line holds for any $s \in \{0, 1, 2, ..., t-1\}$, so that for $s = t-1$, we get the last line.

Steady state implies that $N_{j,t}(x;l) = N_{j,t}(x;l+1)$ for all l, t, j, which leads to $N_j(x;l+1) = N_j(x;l)$ and $p_{l+1} = p_l$. Also, if we have $N_j(x;l+1) = N_j(x;l)$, we get that $p_{l+1} = p_l$, and then $N_{S,t}(x;l) = N_{S,t}(x;l+1) = N_{S,t}(x)$. So a stable total population in both roles yields a stable population in each type.

Suppose $N_S(x;l+1) - N_S(x;l) > 0$, what does this imply for $N_S(x;l+2) - N_S(x;l+1)$? More sellers in period $l+1$, implies a lower finding rate for sellers in that period $\frac{p_{l+1}}{N_S(x;l+1)} < \frac{p_l}{N_S(x;l)}$, so a higher fraction of sellers go on to $l+2$, so the difference between sellers in $l+2$ and sellers in $l+1$ is smaller than it was in the previous period: $N_S(x;l+1) - N_S(x;l) > N_S(x;l+2) - N_S(x;l+1)$. This will continue until the populations converge. The same process occurs for buyers. Therefore, after sufficiently many periods, $N_j(x;l), p_l$ becomes independent of l. So $N_{j,t}(x;l+1)$ also becomes independent of l.

So

$$
\begin{aligned}
N_j(x) = \sum_{t=0}^{\infty} N_{j,t}(x) = \sum_{t=0}^{\infty} \delta N_{j,t-1}(x)(1 - p\frac{1}{N_j(x)}) \\
= N_{j,0}(x) \sum_{t=0}^{\infty} \delta^{t-1}(1 - p\frac{1}{N_j(x)})^{t-1} \\
\equiv N_{j,0}(x) \sum_{t=0}^{\infty} \delta^{t-1}(1 - \omega_j(x))^{t-1} \\
= N_{j,0}(x) \frac{1}{1 - \delta(1 - \omega_j(x))}
\end{aligned}
$$

Where we define $\omega_j(x) = \frac{p}{N_j(x)} = \frac{M(N_S(x), N_B(x))}{N_j(x)}$

QED.

If agents of role j choose to exit at the start of tenure T_j, nothing changes except that the tenure, for the purpose of calculating meeting rates becomes limited by $T_j - 1$. In this case,

$$
\begin{aligned}
N_j(x) = N_{j,0}(x) \sum_{t=0}^{T_j-1} \delta^{t-1}(1 - \omega_j(x))^{t-1} \\
= N_{j,0}(x)[\frac{1 - [\delta(1 - \omega_j(x))]^{T_j-1}}{1 - \delta(1 - \omega_j(x))}]
\end{aligned}
$$

Introducing Market Makers (MMs) who meet with traders at a known rate of ρ and who make offers that are acceptable to every agent they meet, causes tenure t population to be:

$$
\begin{aligned}
N_{j,t}(x; l+1) = \delta(N_{j,t-1}(x; l) - \rho N_{j,t-1}(x; l) - p_l \frac{(1-\rho)N_{j,t-1}(x; l)}{(1-\rho)N_j(x; l)}) \\
= \delta N_{j,t-1}(x; l)(1 - \rho - p_l \frac{1}{N_j(x; l)})
\end{aligned}
$$

To understand the first line, note that $N_{j,t}(x; l+1)$ is made up of all agents who in period l were in $N_{j,t-1}(x; l)$ and haven't exited. $(1-\delta)N_{j,t-1}(x; l)$ have exited exogenously. Of the remaining $\delta N_{j,t-1}(x; l)$, a fraction ρ exited due to trade with MMs. And an additional p_l have exited in all tenures, $\frac{N_{j,t-1}(x;l)}{N_j(x;l)}$ of whom were from $N_{j,t-1}(x; l)$.

Therefore the final population is

$$
\begin{aligned}
N_j(x) &= N_{j,0}(x) \sum_{t=0}^{T_j-1} \delta^{t-1}(1-\rho-p\frac{1}{N_j(x)})^{t-1} \\
&\equiv N_{j,0}(x) \sum_{t=0}^{T_j-1} \delta^{t-1}(1-\rho-\omega_j(x))^{t-1} \\
&= N_{j,0}(x)\frac{1-[\delta(1-\rho-\omega_j(x))]^{T_j-1}}{1-\delta(1-\rho-\omega_j(x))}
\end{aligned}
$$

D dPD/dk

Recall from appendix B that:[22]

$$
r_{B,1,1} - r_{S,1,1} = \frac{2(1-d_{B,1,1}-d_{S,1,1}+d_{B,1,1}d_{S,1,1})}{2-d_{B,1,1}-d_{S,1,1}}
$$

Therefore:

[22]Once again, this is assuming $c_j = 0$, the derivation for $c_j > 0$ is analogous

$$
\begin{aligned}
\frac{\partial}{\partial k}(r_{B,1,1}-r_{S,1,1}) &= \frac{2(-\frac{\partial}{\partial k}d_{B,1,1}-\frac{\partial}{\partial k}d_{S,1,1}+\frac{\partial}{\partial k}d_{B,1,1}d_{S,1,1}+d_{B,1,1}\frac{\partial}{\partial k}d_{S,1,1})}{2-d_{B,1,1}-d_{S,1,1}} \\
&- [-\frac{\partial}{\partial k}d_{B,1,1}-\frac{\partial}{\partial k}d_{S,1,1}]\frac{2(1-d_{B,1,1}-d_{S,1,1}+d_{B,1,1}d_{S,1,1})}{[2-d_{B,1,1}-d_{S,1,1}]^2} \\
&= \frac{2(-\frac{\partial}{\partial k}d_{B,1,1}-\frac{\partial}{\partial k}d_{S,1,1}+\frac{\partial}{\partial k}d_{B,1,1}d_{S,1,1}+d_{B,1,1}\frac{\partial}{\partial k}d_{S,1,1})(2-d_{B,1,1}-d_{S,1,1})}{[2-d_{B,1,1}-d_{S,1,1}]^2} \\
&- [-\frac{\partial}{\partial k}d_{B,1,1}-\frac{\partial}{\partial k}d_{S,1,1}]\frac{2(1-d_{B,1,1}-d_{S,1,1}+d_{B,1,1}d_{S,1,1})}{[2-d_{B,1,1}-d_{S,1,1}]^2} \\
&= \frac{2(\frac{\partial}{\partial k}d_{B,1,1}[d_{S,1,1}-1]+\frac{\partial}{\partial k}d_{S,1,1}[d_{B,1,1}-1])(2-d_{B,1,1}-d_{S,1,1})}{[2-d_{B,1,1}-d_{S,1,1}]^2} \\
&+ \frac{2(1-d_{B,1,1}-d_{S,1,1}+d_{B,1,1}d_{S,1,1})(\frac{\partial}{\partial k}d_{B,1,1}+\frac{\partial}{\partial k}d_{S,1,1})}{[2-d_{B,1,1}-d_{S,1,1}]^2} \\
&= \frac{\partial}{\partial k}d_{B,1,1}\frac{2([d_{S,1,1}-1])(2-d_{B,1,1}-d_{S,1,1})}{[2-d_{B,1,1}-d_{S,1,1}]^2}+\frac{\partial}{\partial k}d_{S,1,1}\frac{2([d_{B,1,1}-1])(2-d_{B,1,1}-d_{S,1,1})}{[2-d_{B,1,1}-d_{S,1,1}]^2} \\
&+ \frac{\partial}{\partial k}d_{S,1,1}\frac{2(1-d_{B,1,1}-d_{S,1,1}+d_{B,1,1}d_{S,1,1})}{[2-d_{B,1,1}-d_{S,1,1}]^2}+\frac{\partial}{\partial k}d_{B,1,1}\frac{2(1-d_{B,1,1}-d_{S,1,1}+d_{B,1,1}d_{S,1,1})}{[2-d_{B,1,1}-d_{S,1,1}]^2} \\
&= \frac{\partial}{\partial k}d_{B,1,1}\frac{2(3d_{S,1,1}-2-d_{B,1,1}d_{S,1,1}+d_{B,1,1}-d_{S,1,1}^2)}{[2-d_{B,1,1}-d_{S,1,1}]^2}+\frac{\partial}{\partial k}d_{B,1,1}\frac{2(1-d_{B,1,1}-d_{S,1,1}+d_{B,1,1}d_{S,1,1})}{[2-d_{B,1,1}-d_{S,1,1}]^2} \\
&+ \frac{\partial}{\partial k}d_{S,1,1}\frac{2(1-d_{B,1,1}-d_{S,1,1}+d_{B,1,1}d_{S,1,1})}{[2-d_{B,1,1}-d_{S,1,1}]^2}+\frac{\partial}{\partial k}d_{S,1,1}\frac{2(3d_{B,1,1}-2-d_{B,1,1}^2-d_{S,1,1}d_{B,1,1}+d_{S,1,1})}{[2-d_{B,1,1}-d_{S,1,1}]^2} \\
&= \frac{\partial}{\partial k}d_{B,1,1}\frac{2(2d_{S,1,1}-1-d_{S,1,1}^2)}{[2-d_{B,1,1}-d_{S,1,1}]^2}+\frac{\partial}{\partial k}d_{S,1,1}\frac{2(2d_{B,1,1}-1-d_{B,1,1}^2)}{[2-d_{B,1,1}-d_{S,1,1}]^2} \\
&= -2\frac{\frac{\partial}{\partial k}d_{S,1,1}(d_{B,1,1}-1)^2+\frac{\partial}{\partial k}d_{B,1,1}(d_{S,1,1}-1)^2}{(2-d_{B,1,1}-d_{S,1,1})^2}
\end{aligned}
$$

E Equilibrium

Definitions and Uniqueness

Define $\Gamma \equiv [\Gamma_1, \Gamma_2, ...]$ to be the set of all perfect Bayesian steady state pure strategy reservation price (candidate) equilibria without full trade in this model. To clarify, these are all the equilibria where the stocks of the various types are fixed over time, where agents do not randomize their offer or response strategies, where response choices are made according to the reservation prices, where all strategies are sequentially rational given beliefs and where beliefs are sequentially rational given strategies and signals.

In every such equilibrium, an agent of role $j = \{B, S\}$ and history h, for all h, makes an offer that does not depend on who she is meeting (because she cannot distinguish between them). This offer is equal to the reservation price of some type $(-j, h')$. To see why, first note that I am assuming that such offers are accepted by such type $(-j, h')$. By the definition of reservation prices, they are also accepted by all types with lower continuation values than type $(-j, h')$. Whenever the offer does not equal the reservation price of any type, there exists an offer that is strictly better for the offerer if it is accepted and has an equal probability of being accepted.

Define $\Omega^{(i)}$ to be the set of all histories such that in equilibrium i, there is a strictly positive probability of meeting an agent with such history. Because matching is random, the following definition is adaquate: $\Omega^{(i)} \equiv \{(j,h)|N^{(i)}_{j,h}/N^{(i)}_j > 0\}$, where $N^{(i)}_{j,h}$ is the number of agents of type (j, h) in equilibrium i and $N^{(i)}_j \equiv \sum_h N^{(i)}_{j,h}$.

In each Γ_i, there exists at least one type (j, h_0) whose equilibrium path offer is such that some agent on the other side rejects it. This offer is equal to the reservation price of type $(-j, h')$. h' is different from the history H which yields the highest continuation value for agents of role $-j$ of all histories in Ω. Formally: $o_{j,h_0} = r_{-j,h'} \neq r_{-j,H}$, where $V^{(i)}_{-j,H} \geq V^{(i)}_{-j,h}, \forall h \in \Omega$. Otherwise, Γ_i is a full trade equilibrium.

If Γ_i is an equilibrium, then $R^{(i)}_{j,h'} \cdot A^{(i)}_{j,h'} \geq 1$, where $R^{(i)}_{j,h'} = \frac{1-V^{(i)}_{-j,h'}-\delta_j\cdot V^{(i)}_{j,(h_0,rej(h'))}}{1-V^{(i)}_{-j,H}-\delta_j\cdot V^{(i)}_{j,(h_0,rej(h'))}} \geq 1$, $A^{(i)}_{j,h'} = Pr^{(i)}_{h_0}(acc|h') < 1$. $R^{(i)}_{j,h'}$ is the benefit of type (j, h_0) from offering $r_{-j,h'}$ (which is the equilibrium offer in Γ_i) divided by the benefit to the same type from deviating to offering $r_{-j,H}$, which ensures the offer is accepted. $V^{(i)}_{j,(h_0,rej(h'))}$ is the continuation value after an agent with history h_0, offers $r_{-j,h'}$ and is rejected. $A^{(i)}_{j,h'}$ is the subjective probability that type (j, h_0) assigns to the event that the offer $r_{-j,h'}$ is accepted in the candidate equilibrium.[23]

I am assuming that $k \uparrow$ causes slower learning and $k \to \infty$ causes no learning. This implies that the difference in valuations between two agents with different histories on the same side declines with k. Thus, the benefit from non deviation diminishes ($\frac{\partial}{\partial k}R^{(i)}_{j,h'} < 0$) and converges to 1 for large enough k ($lim_{k\to\infty}R^{(i)}_{j,h'} = 1$), while the utility cost of making an offer that some types would reject remains strictly positive. To see why $\frac{\partial}{\partial k}R^{(i)}_{j,h'} < 0$ and $lim_{k\to\infty}R^{(i)}_{j,h'} = 1$, note that the only reason agents of role j are different in this model from each other is that their different history has induced them to have different beliefs. k governs the extent to which this differentiation occurs. If $k \to \infty$, the $V^{(i)}_{-j,h} - V^{(i)}_{-j,h'} \to 0$ for any two histories h, h' because when agents have identical beliefs, their continuation values are identical. Similarly, at any level of k, increasing k causes $V^{(i)}_{-j,h_0} - V^{(i)}_{-j,h'}$ to

[23]These terms come from the non deviation condition: $1 - V^{(i)}_{-j,H} \leq (1 - V^{(i)}_{-j,h'}) \cdot Pr^{(i)}_{h_0}(acc|h') + (1 - Pr_{h_0}(acc|h'))\delta_j V^{(i)}_{j,(h_0,rej(h'))}$. $R^{(i)}_{j,h'} \geq 1$ is because otherwise $V^{(i)}_{-j,h'} > V^{(i)}_{-j,H}$, contradicting the assumption that $V^{(i)}_{-j,H} \geq V^{(i)}_{-j,h}, \forall h \in \Omega$.

diminish.

At the same time, as long as an agent (j, h_0) believes there is a positive measure of agents of type $(-j, H)$ (which is assured because H was chosen from Ω), $A^{(i)}_{j,h'} = Pr^{(i)}_{h_0}(acc|h') < 1$.

Therefore there exists $k^{(i)}$ large enough such that $\forall k \geq k^{(i)} \Rightarrow R^{(i)}_{j,h'} \cdot A^{(i)}_{j,h'} < 1 \Rightarrow \Gamma_i$ is not an equilibrium.

Note that I am ignoring types with $N^{(i)}_{j,h}/N^{(i)}_j = 0$. If the type space was continuous (if time was continuous for example), it would be more difficult to prove $A^{(i)}_{j,h'} = Pr^{(i)}_{h_0}(acc|h') < 1$. One would have to rule out equilibria with only infinitesimal measures of agents who reject equilibrium offers. However, the following underlying logic should still guarantee some area of uniqueness: as k becomes arbitrarily large, agents becomes arbitrarily similar. $A^{(i)}_{j,h'}$ is converging with k, but it is converging not to 1 but to the probability of acceptance assigned by an agent with the null history. However close this value is to 1, it should be possible to find k large enough such that $R^{(i)}_{j,h'} \cdot A^{(i)}_{j,h'} < 1$. On the other hand, if $A^{(i)}_{j,h'} = 1$, then the equilibrium in question is indistinguishable from FTE.

Existence of FTE

Now, consider the conditions for existence of FTE: that no type (j, h) prefers to deviate from offering $r_{-j,H}$[24]. In such case, by the definition of $r_{-j,H}$ as the reservation price of the counterpart type with the highest continuation value, no deviation to rejecting such an offer is strictly welfare improving. It remains to consider deviations in the offer. The condition for non deviation in the offer is familiar: $R^{(FTE)}_{j,h'} \cdot A^{(FTE)}_{j,h'} \leq 1$, where $R^{(FTE)}_{j,h'} = \frac{1 - V^{(FTE)}_{-j,h'} - \delta_j \cdot V^{(FTE)}_{j,(h_0, rej(h'))}}{1 - V^{(FTE)}_{-j,H} - \delta_j \cdot V^{(FTE)}_{j,(h_0, rej(h'))}} \geq 1$, $A^{(FTE)}_{j,h'} = Pr^{(FTE)}_{h_0}(acc|h')$. Note that now the required sign is reversed because now the equilibrium strategy is offering $r_{-j,H}$ which is always accepted. To rule out a deviation in the offer, the condition should hold for any pair of histories h_0, h', the first of the offerer and the second of the receiver. Once again, by increasing k, $R^{(FTE)}_{j,h'}$ can be made arbitrarily close to 1, while $A^{(FTE)}_{j,h'}$ remains strictly below 1 because I am considering deviations to offers that are rejected by at least one type that one meets with a strictly positive probability.

So there must exist some threshold $\bar{k} = sup[k^{(i)}, k^*], \forall i \in \Gamma$, so $\forall k \geq \bar{k}$, the above holds for all equilibria in Γ, and in such case FTE is the unique steady state pure strategy reservation price equilibrium.

For the sake of concreteness, when the prior is $U[0, 1]$, $\mu_{-j,1,1} - \mu_{-j,0,0} = \frac{1}{6}$ and $R_{j,t} \cdot A_{j,t} \leq 1$ holds whenever $\delta_S = \delta_B \leq 0.999$. Note that $U[0, 1]$ has a greater variance than any single peaked

[24] As defined above, H is the history that yields the highest continuation value for a participant on side j

distribution over $[0,1]$ so whenever FTE holds for $U[0,1]$, it should also hold for all symmetric single peaked distributions over $[0,1]$.

Partial Trade

In partial trade, some agents make offers that some agents on the other side reject (I refer to such offers as risky). The agents who make the risky offers must be the more optimistic agents (optimism is measured as the level of the continuation value). Therefore, agents who match in their first attempt must make risky offers. If they do not, but rather make an offer acceptable to all types on the other side (I refer to such offers as safe), then agents who match only in their second attempt also make a safe offer. By induction, all agents make safe offers and we are back in full trade.[25]

Suppose there are 4 equilibrium transaction prices: $r_{B,Risky} > r_{B,Safe} > r_{S,Safe} > r_{S,Risky}$

The safe offers satisfy the most optimistic agents. These agents have more than 1 match, and for each match except the last one, have either rejected an offer or their offer had been rejected. Suppose that a failed match (match and getting a risky offer or making a risky offer that is rejected) is always a negative signal (in the sense that it causes the continuation value to decline).[26] In this case, beliefs decline with tenure when measured before matches are announced. Then, the most optimistic agents are those that match on their first tenure before offers are made. Their beliefs determine $r_{B,Safe}, r_{S,Safe}$ as in FTE. Therefore, the effect of k on $r_{B,Safe}, r_{S,Safe}$ is similar to the effect of k on $r_{B,1,1}, r_{S,1,1}$ in FTE.

The effect of k risky offers are more difficult to analyze because they are determined both by 3 separate effects.

1. The effect on the beliefs of the type that is targeted by risky offers to be indifferent to that offer (so that all those weakly more pessimistic accept). These agents are pessimist, otherwise it does not pay to risk rejection in targeting them. Therefore, a higher k makes them more optimistic, so sellers demand more and buyers are willing to pay less $\downarrow r_{B,Risky}, \uparrow r_{S,Risky}$, which makes the offerers of risky offers worse off.

[25] The fact that tenure 1 agents make a risky offer implies that unless the exit rate is very low, the probability of receiving a risky offer conditional on matching is non negligible. This is because tenure 1 agents are the most numerous agents, so as long as population stocks are finite, they are a strictly positive fraction of agents. The rejectors of risky offers must be the more optimistic agents on the other side, so they must include agents that match on their first tenure. So again, unless the exit rate is very low, the probability of rejection is non negligible.

[26] As long as a failed match is a positive signal, the beliefs of some agents increase with tenure. if a failed match is never a bad signal, then the most optimistic agents have an infinite history and an infinite number of matches and very high beliefs. Also, there is a negligible amount of them, so targeting them is never an equilibrium. So in this case, there should not exist an equilibrium that features a safe offer. One possibility is a semi-pooling equilibrium where all types make a risky offer, in which case, this can be seen as a special case of the case discussed next.

2. The effect on the beliefs of the types that make risky offers which may affect which type is targeted. At least some of the makers of risky offers are optimistic relative to the prior. Their beliefs decline with a higher k, which makes them believe the matching rate is higher on the other side, so that their offers are less likely to be accepted, which makes them make offers that are acceptable to more agents, so the offerer sellers request less from the buyer they meet $\downarrow r_{B,Risky}$ to be paid less and the offerer buyers bid more $\uparrow r_{S,Risky}$.[27] However, this effect is attenuated or even reversed if enough of the offerers are pessimistic relative to the prior (though they should still be optimistic relative to the makers of safe offers)

3. The effect on the set of types that make risk offers. Suppose $(j, h_{j,0,Risky})$ is the most pessimistic type that makes the risky offer while $(j, h_{j,1,Safe})$ is the most optimistic type that makes the safe offer. Note that $V_{j,h_{j,0,Risky}} > V_{j,h_{j,1,Safe}}$. The effect of k on changing these types depends on whether or not these types are more or less optimistic than the prior. The options are:

   (a) $V_{j,\pi_0} > V_{j,h_{j,0,Risky}} > V_{j,h_{j,1,Safe}}$, where π_0 is a history that only includes the public signal. In this case, assuming $|\pi_0 - \pi^*| \leq \overline{\pi}$ (with appropriate adjustment of A3), k makes both more optimistic so it may shift $h_{j,1,Safe}$ to make a risky offer. This will lowering the expected utility from a match on the other side, which tends to cause earlier quitting. Assuming a symmetric setting and equilibrium, the effect will occur on both sides (otherwise, the lower continuation value on the other side may cause the first side to become even more optimistic, magnifying the effect of k). In the distribution of transaction prices, this will increases the density on the two extreme transaction prices, so price dispersion tends to increase.

   (b) $V_{j,h_{j,0,Risky}} > V_{j,h_{j,1,Safe}} > V_{j,\pi_0}$. In this case, they are both optimistic, so under the above conditions, a higher k makes them less optimistic, and may shift $h_{j,0,Risky}$ to making a safe offer. The offers are the mirror image of 1. In this case, all three effects tend in the same direction so that $\downarrow r_{B,Risky}, \uparrow r_{S,Risky}$ which reduces price range and tends to reduce price dispersion. This has the consequence of increasing the continuation values of the most pessimistic agents, which is added to the effect on k on the rate of learning which also increases their continuation values. Both tend towards postponing quitting, the consequences of which can be analyzed using the framework developed for FTE.

   (c) $V_{j,h_{j,0,Risky}} > V_{j,\pi_0} > V_{j,h_{j,1,Safe}}$. Here, a higher k can have either of the two above effects.

[27] Note that this is not actually an effect on the reservation prices of specific types, but changes in the targeted types.