London School of Economics and Political Science

Essays in Market Microstructure

Francesco Palazzo

A thesis submitted to the Department of Economics of the London School of Economics for the degree of Doctor of Philosophy, London, April 2015.

Declaration

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it).

I warrant that this authorisation does not, to the best of my belief, infringe the rights of any third party.

I declare that my thesis consists of 35,157 words.

Francesco Palazzo

Statement of conjoint work

Chapter 3, "Learning and Price Dynamics in Durable Goods Market", was jointly co-authored with Min Zhang. I contributed a minimum of 50% of this work.

General decisions about the direction of research, and the proofs of the main results, were made equally between the authors.

Francesco Palazzo *Min Zhang*

Acknowledgements

This thesis would not have been possible without the support of many people. I am very grateful to my advisors Dimitri Vayanos and Francesco Nava, for their guidance and encouragement. A special thanks goes to Andrea Prat, for his constant support and interest in my work, and for his generous invitation to Columbia University.

My work has immensely benefitted from long and fruitful discussions with my fellow PhD students. I am extremely grateful to Min Zhang and Jason R. Donaldson; they have been great co-authors and invaluable friends. Special thanks go to Markus Riegler, Ana McDowall, Oliver Pardo Reinoso, Can Celiktemur, Sebastian Kodritsch, and the other participants of the LSE Economic Theory seminar series.

I would like to thank the Bank of Italy for financial support during my PhD studies, and for a stimulating work environment now. Anatoli Segura and Taneli Mäkinen provided very valuable feedback on the work contained in this thesis.

Though there are more people who deserve thanks than I can name here, I would like to particularly thank Marco Bevilacqua, Diletta Drago, Giuseppe Calì, and Mara Giua, for sharing my enthusiastic as well as my tough moments.

Lastly, I cannot thank enough my parents for their support. Without their example I would not know what intellectual curiosity means.

Abstract

This dissertation contains three theoretical essays on the functioning and the organization of over the counter markets.

The first paper, “Is Time Enough to Alleviate Adverse Selection?,” considers a dynamic adverse selection model in which sellers pay a search cost to find a new buyer. I uncover a relationship between adverse selection and the magnitude of search costs. Interestingly, small search costs may increase the severity of the adverse selection problem, ultimately leading to a lemons market. A market design intervention may mitigate adverse selection and promote full market participation. Conditional upon an adequate level of information disclosure, a per period market participation tax, coupled with a final rebate once a seller trades, introduces a credible signalling device.

The second paper, “Peer Monitoring Incentives via Central Clearing Counterparties,” studies how the novel introduction of mandatory clearing for over the counter financial assets may affect dealers’ incentives to monitor each other’s. The design of the loss allocation rules is crucial. To maximize peer monitoring incentives, a higher share of losses should be payed by surviving members with a greater trade exposure to the defaulting dealer. In practice, this mechanism can be implemented through variation margin haircutting. If all members should contribute, equilibrium outcomes may be inferior to what can be achieved without clearing.

The third paper, “Learning and Price Dynamics in Durable Goods Markets,” is joint work with Min Zhang. We set up a dynamic model with two key features: first, agents enjoy heterogeneous use values, and later resell the good; second, prices do not incorporate all available information. Informational frictions slow down learning, and affect price movements asymmetrically in high and low aggregate demand states. Learning and the resale motive are the predominant force for durable goods with short resale horizons, slow time-varying aggregate demand, and similar use values among buyers.

Contents

List of Figures

Chapter 1

Is Time Enough to Alleviate Adverse Selection?

In the dynamic adverse selection literature, postponing trade is a signal of quality: low types trade early, high types demand higher prices and trade afterwards. I reconsider this result in a model in which market outcomes depend on the interaction among (i) the magnitude of sellers' search costs, (ii) the precision of buyers' signals on product quality, and (iii) buyers' ability to observe how long a seller has been on the market. The lemons problem is more severe, ceteris paribus, when search costs are small. A low type sets a high price because it is cheap to wait until a buyer wrongly perceives his product as good. However, the market breaks down if high quality products are not numerous enough. More precise signals and time on market observability mitigate the likelihood of a lemons market. A well designed mechanism may induce all sellers to trade: sellers should pay a per period market participation tax, and receive a rebate after trading.

1.1 Introduction

Information asymmetry is a pervasive feature of real-world markets. Financial securities, real estate, electronics and secondhand vehicles are just a few examples. One side of the market—usually buyers—lacks information or experience to ascertain the true quality of a specific good. Since Akerlof's (1970) seminal paper, it is well known that in a static model 'lemons' may force high quality products out of the market.

A growing literature has been reconsidering the adverse selection problem in a dynamic environment.[1] A key feature of these models is the use of time as a signalling device so that every type of seller eventually trades. In equilibrium, low quality sellers trade early while high quality sellers demand higher prices and trade afterwards. I refer to this economic mechanism as *inter-temporal separation* (henceforth, ITS).[2] Postponing trade signals good quality, and opens up the opportunity to sell at a higher price. Although this inefficient delay among high quality sellers causes some welfare loss, eventually every seller trades.

Despite ITS theoretical importance, empirical findings seem at odds with its mechanism. Tucker et al. (2013) point out that real estate sellers who have been on the market longer trade at lower prices, when buyers can credibly observe how long a seller has been on the market. Trading late does not seem to strengthen reputation, but instead it is perceived to signal lower quality. Furthermore, ITS fails to address Jin and Kato's (2007) evidence on the relationship between adverse selection and market segmentation. They show that the lemons problem may induce different product qualities to separate between online and offline markets rather than over time. Specifically, products offered online are more likely to be of low quality, unless certified by a professional third party, while higher quality products are usually sold through the retail channel.[3] Lastly, Lewis (2011) points out that greater information disclosure on eBay motors increases sellers' chances of trading as well as final prices; nevertheless, inter-temporal separation does not explain this piece of evidence, as its mechanism works irrespective of the existence of informative signals for buyers. Together, the empirical findings of these studies suggest to reconsider—at least for some markets—how time may mitigate adverse selection in real-world markets.

This paper proposes a model to explain why these empirical patterns emerge, and how a market designer may induce high quality sellers to participate in an otherwise

[1] A non-exhaustive list includes Janssen and Roy (2002), Blouin (2003), Camargo and Lester (2014), Fuchs and Skrzypacz (2013, 2014), and Kaya and Kim (2014).

[2] A key assumption underlying the inter-temporal separating mechanism is the ability of buyers to infer how long a seller has been waiting on the market. They can either observe hard evidence of time on market or there is common knowledge on the initial date of the game.

[3] Analogously, Dewan and Hsu (2004) find that identical stamps trade at a 10–15 percent discount on eBay compared to a specialty stamps auction with lower quality uncertainty.

lemons market. In my setup, market outcomes depend on the interaction among three main features of the dynamic sale problem. First, every seller incurs a search cost to find a new buyer and make a price offer; second, once matched with a seller, the buyer receives a binary informative signal on the specific good offered on sale; lastly, buyers may or may not have credible information on sellers' time on market. I use the acronyms TMO and TMN to denote when time on market is observable or not. I characterize equilibria for TMO and TMN separately to understand whether revealing this information improves final allocations.

My main result is that markets with *low search costs* suffer from a more severe adverse selection problem. Intuitively, a low quality seller pools on the same price offered by high quality sellers because it is cheap to wait until a buyer will wrongly perceive his product as good. Due to this imitating strategy, high quality sellers may decide to stay out of such markets or, at best, to participate for a limited time only. This result is in contrast with the ITS logic, which maintains that sellers separate over time for each discount factor—a measure of search frictions—even if trading may require them to spend, on average, a long time on the market.

My main departure from models featuring ITS is an alternative assumption about the delay cost that sellers incur to find a new trade opportunity. This difference is not just a technical issue about preference specifications, as it captures two alternative economic ideas on the nature of these costs. In models with ITS, postponing trade imposes a delay cost via time discounting, reducing the present value of a positive expected payoff. Trading late is costly, but market participation always provides a non-negative payoff. In contrast, in my model postponing trade imposes a per period cost in the form of an additive utility loss and, for simplicity, there is no discounting.[4,5] Sellers may stay out of the market if they expect to incur a considerable cumulative cost before trading.[6] This *per period* cost could take many plausible forms: for example, search costs, market participation fees, maintenance costs, or in a financial market it could capture the cost of carrying an open position. For simplicity, I refer to this utility loss as a 'search cost' but alternative interpretations are possible depending on the market under consideration. The main point is to have an environment in which postponing trade for a sufficiently long time dissipates all gains from trade.

[4] I take a neutral stance and consider an equal cost for each type of seller. To the best of my knowledge, there is no particular economic reason to assume different costs for different seller' types. However, the main economic mechanism presented in this paper would still be valid if high quality sellers did not enjoy a significant cost advantage.

[5] To emphasize the different economic mechanism at work, I assume no payoff discounting. Atakan (2006) and Lauermann and Wolinsky (2013) assume the same preference specification. Analogous results would hold if discounting has an order of magnitude sufficiently small relative to search costs; see Example 1.8.1.

[6] From a technical standpoint, ITS relies on discounting of an instantenous payoff. This preference specification guarantees two essential properties: first, a strict single crossing condition with respect to time, and, second, perpetual market participation as waiting costs cannot lead to a negative expected payoff. In contrast, in my setup, finding a new trade opportunity imposes an additive and symmetric (w.r.t to sellers' types) cost. As a result, cumulative delay costs may be larger than total gains from trade.

In my model, separation is possible when a low quality good trades at a low price that buyers always accept, whereas a high quality one trades at a higher price to the first buyer who receives a high signal on its quality. In other words, separation is based on the difference in the expected search costs of pursuing a high price strategy by the two types of sellers. More precise signals for buyers or higher search costs for sellers increase this difference. For a given search cost, higher signal precision decreases the expected cost for high quality sellers but increases the cost for low quality types. For a given signal precision, a higher search cost increases the expected cost for both types of sellers but low quality ones suffer a larger loss. Separation is possible only if high quality sellers enjoy a sufficient advantage in terms of expected search costs, i.e. if signal precision *and* search costs are sufficiently high.[7] This separating equilibrium does not depend on whether the market exhibits TMO or TMN.

When buyers' signals are not very informative and/or search costs are too small,[8] the difference in expected search costs of a high price strategy does not prevent low quality sellers from pooling on a high price whenever possible. In turn, only two outcomes exist: either the market includes only lemons, or all sellers participate and post the same 'high' price, accepted only after a positive signal. A market breakdown is inevitable when the share of high quality products among the entrant sellers is below a certain threshold value, which depends on signal precision and time on market observability. Specifically, a lemons market is more likely to emerge when buyers do not have sufficiently informative signals, in line with Lewis's (2011) findings. However, my model points out a complementarity between signal precision and search costs. In this respect, the potential for a market breakdown when search costs are small may provide a rationale for Jin and Kato's (2007) evidence on the likely exclusion of uncertified high quality sport cards from eBay.[9] Online markets have almost eliminated search costs and—according to my model predictions—low types pretend to offer a high quality good, because it is cheap to wait for a buyer who receives a positive signal and accepts a high price. However, in equilibrium, prices do not convey any information on the underlying quality of the good and buyers are skeptical to accept a high price. In turn, owners of high quality goods may prefer to avoid online platforms. For example, they could either sell through an a certified intermediary, or offer a costly contractual device (for instance warranty) to alleviate adverse selection. I do not model these alternatives, but they are implicitly captured by the type-dependent reservation

[7] Obviously, search costs should not be *too* high; otherwise no seller would participate in the market.

[8] A literature on e-commerce focuses on competition and price dispersion (Clay et al. (2001) and Baye et al. (2004)) or on market structure (Goldmanis et al. (2010)). However, to the best of my knowledge, no existing paper specifically discusses how low search costs interact with the adverse selection problem.

[9] However, as argued by Dellarocas (2005), eBay has been successful in solving asymmetric information problems with adequate mechanism design and feedback. I consider my model more closely related to other less sophisticated online platforms such as Craigslist.com, Gumtree.com or Pianomart.com, where anonymous sellers place simple ads and buyers can individually contact the seller.

value that every seller enjoys. My focus is the relationship between delay costs and adverse selection; hence, I exclude other potential signalling devices.

In a pooling equilibrium—when search costs are low—information on sellers' time on the market is crucial. It determines buyers' prior beliefs about the likelihood of matching with a high quality seller. If time on market is observable, buyers' prior probability of receiving a high quality good from a seller κ periods old, say, is equal to the share of high quality sellers in cohort κ. If time on market is not observable, buyers hold a single prior equal to the overall share of high quality sellers in the market. The two alternative assumptions lead to different price dynamics. Because posted prices are accepted only after a high signal, under TMO the longer a seller has been on the market the lower his price offers will be: high quality goods sell more rapidly—as they are more likely to receive a high signal—and buyers realize that older sellers are more likely to offer a low quality good.[10] If a seller has been on the market too long, buyers would only accept prices below the reservation value of high quality sellers who, in turn, prefer to exit the market. The model predictions of a decreasing price path and an eventual market drop out are consistent with the empirical patterns reported in Tucker et al. (2013) and Hendel et al. (2009).[11]

When time on market is not observable, everyone offers the same price and no seller drops out once he initially decides to participate in the market. However, high quality sellers are less likely to participate compared to the TMO case. Interestingly, under TMN neither the dynamic dimension, nor the existence of private informative signals for buyers, improves market allocations relative to a static model with uninformed buyers. In markets with more precise signals, low quality sellers wait longer until a buyer receives a positive signal, and the pool of sellers on the market includes a larger share of low quality goods relative to the cohort of entrant sellers. This negative effect on the prior probability to receive a high quality good perfectly cancels out the positive effect of a more precise signal. In other words, under TMN, increasing signal precision is self-defeating as it worsens the average market quality which, in turn, determines buyers' prior expectations of receive a high quality good. Thus, the ability to observe time on market may improve welfare when the number of high quality goods is low.

In light of these negative results, I perform in section 1.6 a market design exercise for the limit case of zero search costs. I analyze whether a system of transfers conditional on market participation and trade may alleviate adverse selection and promote

[10] Taylor (1999) is the first paper to exploit this social learning mechanism in the context of a two-period adverse selection model (see section 2.2).

[11] Hendel et al. (2009) document that some real estate sellers in Madison, WI decide to switch to a realtor after some time spent on a for-sale-by-owner website, an online platform with a publicly observable posting day. Although they do not discuss how their results relate to the adverse selection problem, my model predictions under TMO match sellers' decision to abandon the online platform after some time.

full market participation. I focus on mechanisms that satisfy a series of properties: budget balance, informational efficiency of prices, and interim individual rationality. The efficient market design intervention achieves separation through a constant market participation tax, and it relaxes sellers' individual rationality constraint through a final rebate conditional on trade. A low quality seller does not find it profitable to post a high price because, on average, he is less likely to find a buyer who receives a high signal; if he pursued a high price strategy, he would pay, on average, a cumulative amount of market participation taxes that would make imitation unprofitable. In terms of incentive compatibility constraints, the market participation tax is analogous to a per period search cost. Nevertheless, the former is not a waste of economic resources, and it can be partially recouped through a rebate, relaxing sellers' market participation constraints. Although time on market observability plays a relevant role in all pooling equilibria, it does *not* affect the efficient mechanism. Taxes and rebates are inversely proportional to buyers' signal precision, but they do not depend on sellers' time on market, although in principle they could. This efficient market design intervention achieves full market participation in a large set of economies, but it is not successful when buyers' signals are close to being uninformative.

From a technical standpoint, the model is a dynamic signalling game since in every period the informed party—sellers—decides to post the price at which they are willing to trade. In previous non-stationary models, the bargaining protocol either assumes exogenous prices[12] or buyers make take-it-or-leave-it offers.[13] As in the Diamond's (1971) paradox, the latter protocol implicitly fixes the price at which high quality sellers trade—equal to their exogenous reservation value—and only the price accepted by low quality sellers is determined endogenously. In my setup, this bargaining solution leads to a hold-up problem: high quality sellers would not pay a search cost to trade at their reservation value. I assign all bargaining power to sellers to improve their chances of participating in the market. Equilibrium characterization is challenging because I have to take into account—for every cohort of sellers—an endogenous behavioural strategy (possibly mixed) for each type of seller.

In the next section I discuss the related literature. Section 3 presents the model setup. Section 4 characterizes the equilibria when search costs are close to zero. Section 5 discusses the welfare properties of equilibria. Section 6 derives the efficient intervention. Section 7 concludes. All proofs are in Appendices A and B.

[12] See Wolinsky (1990), Blouin (2003) and Camargo and Lester (2014).

[13] See Moreno and Wooders (2010, 2014), Kim (2014) and Kaya and Kim (2014).

1.2 Related literature

This paper is mainly related to the theoretical literature on dynamic adverse selection in decentralized markets.[14] Two different types of goods coexist in the same market, but product quality is sellers' private information. The literature mainly considers non-stationary equilibria, as the market starts at an initial date and strategies depend on time.[15] Analogously, the TMO case in this paper leads to a non-stationary equilibrium, since sellers' strategies generally depend on previous time on market. In my setup, time on market coincides with the number of previous matches with buyers.[16,17]

Blouin (2003), Camargo and Lester (2014), and Moreno and Wooders (2014) characterize non-stationary equilibria in infinite horizon games.[18] Inter-temporal separation allows all sellers to trade over time. The main differences among these papers have to do with the division of trade surplus and are partly driven by alternative bargaining protocols. The former two papers adopt the exogenous price bargaining of Wolinsky (1990), while the latter assume buyers make take-it-or-leave-it offers. Kaya and Kim (2014) construct a model in which buyers receive private informative signals and make offers to sellers. In their setup, prices and beliefs converge to a steady state, and the transition depends on the initial probability of trading with a high quality seller: if it is high, prices and beliefs move downward as in Taylor (1999); if it is low, ITS kicks in and allows sellers to separate. As discussed in section 2.3, I assume sellers make take-it-or-leave-it offers. In contrast to what happens in Moreno and Wooders (2010, 2014) or Kaya and Kim (2014), a hold-up problem would arise in my setup if H-sellers could only trade at their reservation value v_H. This is not a concern in models that use discounting of an instantaneous payoff: in equilibrium high quality sellers discount a zero payoff as they trade at their reservation value. My TMO setup is a non-stationary dynamic signalling game with endogenous prices, and the main challenges arise because H-sellers' posted prices are generally non-constant over time.

[14] The latter term defines a class of models that depart from the classic Walrasian price formation paradigm to explicitly model the bilateral interaction between buyers and sellers. A non-exhaustive list of previous papers on decentralized markets with complete product information includes Diamond (1971), Rubinstein and Wolinsky (1985), Gale (1986a,b), Duffie et al. (2005, 2007), Vayanos and Weill (2008), and Lagos and Rocheteau (2009). Wolinsky (1990) considers a decentralized market with asymmetric information on the common quality of all units. Serrano and Yosha (1993), Blouin and Serrano (2001), and Duffie et al. (2009, 2014) provide other contributions to this literature.

[15] Daley and Green (2012) analyze a dynamic setting in which buyers receive public information on the asset value at random arrival times. Buyers may enter a waiting period: if good news arrives confidence is restored and the market reopens; otherwise, there is a partial sell off of low value assets.

[16] I prefer to use the expression 'time on market' for lexical convenience.

[17] Kim (2014) shows that when market frictions are small (small discount rate), observing only time on market is welfare-improving relative to public information of previous matches. This result stems from the fact that staying on the market strengthens reputation; in this respect, information on previous matches conveys a more precise signal than time on market. As a consequence, sellers tend to delay trade as they reject price offers more often. However, my paper points out why this ITS mechanism might not work, and it shows that—when search frictions are small—a longer stay on the market is interpreted as a negative indicator of quality. Therefore, in my setup the welfare comparison between the two regimes could be reversed.

[18] Janssen and Roy (2002) present the ITS mechanism in a model of dynamic centralized competitive markets.

My paper is also related to a new strand of literature on optimal market intervention for lemons markets. Fuchs and Skrzypacz (2013) study how to minimize delay costs through the optimal design of market openings. Fuchs and Skrzypacz (2014) consider government interventions through taxes and subsidies.[19] Their Pareto improving budget balanced policy suggests a short tax exempt trading window followed by a short lived period of positive taxes; sellers trade immediately and after the tax goes back to zero. My efficient intervention also points out the need to subsidize initial trade, but it prescribes a constant market participation tax thereafter (see section 1.6). Fuchs and Skrzypacz's (2014) short-lived taxation policy would not be effective in my setup by the same logic that excludes ITS.

My results are also related to the literature on sequential trading between a long-lived seller and a sequence of short-lived buyers. Taylor (1999) considers a two-period model in which a single informed seller posts prices under different price observability assumptions.[20] His paper was the first to point out the negative informational externality that affects older cohorts of sellers when buyers observe private informative signals. Lastly, Lauermann and Wolinsky (2013) consider a sequential search model with TMN, informative private signals for buyers, and additive search costs. They show the existence of a search friction that reduces price informativeness compared to a common auction environment. They consider buyers who receive signals sampled from a continuous distribution—possibly of unbounded precision—while I use a simple symmetric binary signal of bounded precision. My choice is motivated by tractability concerns, especially for the non-stationary equilibria under TMO. Moreover, I focus on allocative efficiency and market exclusion, while they analyze informational efficiency.

1.3 Model

This section presents the model setup and discusses how the main assumptions relate to the research questions.

1.3.1 Model setup

Consider a decentralized market where trade is possible only in bilateral transactions between one buyer and one seller.[21] Each seller is endowed with a single indivisible

[19] Their paper differs from Philippon and Skreta (2012) and Tirole (2012) because the latter consider a governement intervention in the presence of a static competitive private market.

[20] Hörner and Vieille (2009) and Fuchs et al. (2014) also study the effect of price history observability in models in which buyers have no informative signals.

[21] As a convention, throughout the paper I refer to the seller as "he" and to the buyer as "she".

good of high (H) or low (L) quality. A seller knows the quality of his product but nobody else can observe it. Let θ_λ and v_λ be buyers and sellers' valuation, respectively, for a product of quality $\lambda \in \{H, L\}$, and assume $\theta_H > v_H > \theta_L > v_L$. I use the terms H-sellers and L-sellers to refer to sellers with goods of high and low quality.

Time $t \in \{.., -1, 0, 1, ..\}$ is discrete and in each period a set of sellers μ_t of unit mass is born. Only a fraction $q^0 \in (0,1)$, independent of t, of newly born sellers owns a high quality good. Sellers are long lived and they can participate in the market until they trade or exit. Buyers live for a single period and they always outnumber sellers.

I denote the set of sellers participating in the market at time t as S_t, while $S_t^\kappa \subset S_t$ is the set of sellers who have been participating in the market for $\kappa \in \mathbb{N}_0$ previous periods; similarly, $S_{\lambda,t}^\kappa \subset S_t^\kappa$ is the subset of sellers of type λ in S_t^κ.[22] Sellers pay a search cost c to participate in the market and match with a buyer. Buyers match uniformly at random with sellers and have no search cost. For simplicity, they have no opportunity to buy a good and re-sell it on the market. All players are risk-neutral and have quasi-linear utilities with respect to monetary transfers. Sellers do not discount future payoffs.

Buyers and sellers trade according to a simple mechanism. Each seller $i \in S_{t-1} \cup \mu_t$ who has *not* traded at time $t-1$ takes an action $a_{S,i} \in A_S = \{\{D\}, \mathbb{R}_+\}$, where D denotes the decision to irreversibly drop out of the market, and $p \in \mathbb{R}_+$ is the posted price at which he commits to sell the good in period t. If $a_{S,i} = D$, seller i is not matched with a buyer and does not pay the cost c; however, he has no future possibility of participating in the market. If $a_{S,i} = p$, seller i pays c and gets matched with a buyer. A particular history for seller $i \in S^\kappa$ is indicated with $h_i^\kappa = (h_i^{\kappa-1}, a_{S,i}^\kappa \times a_{B,i}^\kappa)$ (with $h_i^{-1} = \varnothing$) and H^κ is the set of all possible histories.

Let $Z_i^\kappa \subset H_i^\kappa$ denote the set of terminal histories for seller i after κ previous periods (with $Z_i = \bigcup_{\kappa \in \mathbb{N}_0} Z_i^\kappa$). If $h_i^\kappa \in Z_i^\kappa$ seller i exits the market after κ previous periods in the market and he cannot choose any further action, i.e. $A_{S,i}^j = \varnothing,\ j \geq \kappa+1$. Let $Z_i^\kappa(D) \subset Z_i^\kappa$ include all terminal histories in which seller i drops out of the market after κ periods; similarly, $Z_i^\kappa(p) \subset Z_i^\kappa$ denotes the set of histories in which seller i trades at price p after κ previous periods in the market. The final payoff to seller $i \in S_{\lambda,t}^\kappa$ in $z \in Z_i^\kappa$ is

$$\tilde{u}_\lambda(z) = \begin{cases} -\kappa c & \text{if } z \in Z_i^\kappa(D) \\ p - v_\lambda - (\kappa+1)c & \text{if } z \in Z_i^\kappa(p) \end{cases}$$

Once matched with seller i, a buyer receives a private signal $\xi \in \{H, L\}$ on his

[22] To simplify exposition, I slightly abuse notation using S, S^κ, S_λ^κ to denote both the set or the measure of sellers in these sets.

product quality, but she cannot observe his previous price history.[23] Buyers' signals have precision $\gamma \in (\frac{1}{2}, 1)$, i.e. $\mathbb{P}_H(\xi = H) = \mathbb{P}_L(\xi = L) = \gamma$. For a given vector $(\theta_H, v_H, \theta_L, v_L)$, I parametrize a specific economy $\mathscr{E}(\gamma, q^0)$ by signal precision γ and newly born measure q^0 of H-sellers.

I consider two different setups for publicly available information. If time on market is observable (TMO), a buyer observes how long a seller has been participating in the market; i.e. it is common knowledge whether $i \in S_t^\kappa$ for some $\kappa \in \mathbb{N}_0$. In contrast, if time on market is *not* observable (TMN) no buyer can observe this information.

When time on market is observable, a buyer's information set $\mathscr{I}_B(p, \kappa, \xi)$ includes the seller's offer p, his previous κ periods in the market, and the buyer's signal ξ. If time on market is not observable, it only includes p and ξ (i.e. $\mathscr{I}_B(p, \xi)$). Given her information set, a matched buyer takes an action $a_{B,i} \in A_B = \{A, R\}$, where A denotes acceptance and R rejection of the seller i price offer. If she accepts offer p, trade occurs and they leave the market; if she rejects, no exchange takes place and seller i moves to period $t+1$.

In this paper, I only consider stationary and symmetric equilibria of the game. Players' equilibrium strategies do not depend on time t, but only on seller's type λ, cohort κ and history $h_i^{\kappa-1}$. In this class of equilibria, the mass of sellers S_t, S_t^κ and $S_{\lambda,t}^\kappa$ is constant over time—i.e. $S_t = S$, $S_t^\kappa = S^\kappa$ and $S_{\lambda,t}^\kappa = S_\lambda^\kappa$ for every $\kappa \in \mathbb{N}_0$ and $t \in \mathbb{Z}$—and I omit the subscript t in the remainder of this paper. I denote with σ a strategy profile and with π a belief system. A strategy profile σ and a belief system π form an assessment (σ, π). I use σ_{-i} and π_{-i} to indicate the strategy profile and the belief system of any agent other than i.

Let $q^\kappa = \mathbb{P}(\theta_H | S^\kappa)$ be the prior probability under uniform random matching that a seller in S^κ offers a high quality good. On the equilibrium path, a buyer incorporates her private signal into the publicly available information according to Bayes' rule.

Definition 1.3.1 *A assessment (σ, π) is an equilibrium of the game if it is a weak Perfect Bayesian Equilibrium (PBE) with the following restrictions:*

1. ***Stationarity:*** *buyers and sellers in $S_{\lambda,t}^\kappa$ and $S_{\lambda,t'}^\kappa$ play identical strategies $\forall t \neq t'$.*
2. ***Symmetry:*** *if sellers $i, j \in S_\lambda^\kappa$ have $h_i^{\kappa-1} = h_j^{\kappa-1}$, they play the same strategy.*
3. ***Pure strategies:*** *buyers only play pure behavioural strategies.*

I impose a few restrictions on the notion of weak PBE.[24] First, strategies do not depend on time t for otherwise identical sellers. Second, sellers' strategies can differ

[23] This assumption significantly simplifies the set of possible equilibria. See Taylor (1999), Hörner and Vieille (2009) and Fuchs et al. (2014) for models that consider equilibria with price history observability.

[24] See Definition 3.3.1 in Appendix A for a formal definition.

only with respect to seller' types $\lambda \in \{H, L\}$, cohorts S^{κ}_{λ} and previous history $h_i^{\kappa-1}$. The third condition restricts buyers—conditional on an information set—to play pure strategies,[25] but sellers can use mixed strategies.

In Propositions 1.4.3 and 1.4.4, I use the undefeated equilibrium refinement introduced by Mailath et al. (1993).[26] I adopt this refinement for two simple purposes: (i) to rule out the self-fulfilling PBE in which only L-sellers trade because buyers believe only L-sellers participate, whenever there exists another PBE in which H-sellers participate in the market; and (ii) to select among the set of pooling PBE the one with the highest possible prices, i.e. $p^{\kappa} = \mathbb{E}_{\pi_B}[\theta | \mathscr{I}_B(p^{\kappa}, \cdot, H)]$. Therefore, it is *not* used to rule out separating or semi-separating strategy profiles, differently from what happens in the Spence (1973) model.[27] Indeed, Lemma 1.4.1—the main characterization result for small c—does not rely on the undefeated refinement to show that the only admissible behavioural strategies have both types of sellers in S^{κ} pool on the same price. De facto the undefeated refinement is not restrictive, but instead is a conservative choice to illustrate that a market breakdown is still possible under a dynamic setup. Indeed, this refinement selects the equilibrium in which H-sellers' market participation constraint is satisfied for the lowest possible q^0.[28]

For simplicity, I do not specify out-of-equilibrium beliefs in the proposition statements. The main result on the admissible equilibrium strategies—Lemma 1.4.1—rules out other strategies without relying on any specific out-of-equilibrium belief. The resulting admissible equilibria only require buyers to hold sufficiently pessimistic beliefs out of the equilibrium path.

1.3.2 Discussion of the assumptions

I briefly discuss the main model assumptions and how they relate to the literature.

Additive search cost c. My main departure from the previous literature—with the notable exception of Lauermann and Wolinsky (2013)—is the introduction of a search cost. It can be alternatively interpreted as an additive and symmetric specification of delay costs. This preference specification has two main properties: (i) sellers stay out of the market if they expect to trade after a long time, since the cumulative amount of search costs would be larger than their total gains from trade; and (ii) all sellers suffer the same utility loss if they postpone trade.[29] Additive delay costs are not the only

[25] Buyers' strategies depend on the information set, and this restriction does not prevent strategies to depend on signal ξ despite the latter may not change the belief π_B. To understand the logic of this restriction see footnote 35.

[26] See Definition 1.8.2 in the Appendix for a formal definition.

[27] See Mailath et al. (1993) for a discussion.

[28] Roughly speaking, H-sellers in S^{κ} participate only if $\mathbb{E}_{\pi_B}[\theta | \mathscr{I}_B(p^{\kappa}, \cdot, H)] \geq p^{\kappa} \geq v_H$, and, in a pooling equilibrium, this expectation is strictly increasing in π_B, which, in turn, is weakly increasing in q^0.

[29] Atakan (2006) highlights the role of asymmetric delay costs in a model of assortative matching.

preference specification with these properties,[30] but I prefer an additive cost because it can be easily interpreted as a search cost. Moreover, it provides a natural starting point for my market design exercise in which transfers enter utility additively.

Bargaining protocol. All previous papers in the literature adopt a specific bargaining protocol to make the model tractable. Blouin (2003) and Camargo and Lester (2014) adopt the exogenous price bargaining protocol first introduced by Wolinsky (1990). Moreno and Wooders (2010) and Kaya and Kim (2014) assume that buyers make take-it or leave-it offers; Lauermann and Wolinsky (2013) use a random proposals bargaining model to avoid dealing with out of equilibrium beliefs. In this paper, sellers—the fully informed party—make take-it-or-leave-it offers to buyers. In addition to being a realistic assumption in many real-world markets, this trade protocol gives full bargaining power to sellers. Since high quality sellers may stay out of the market, this assumption seems a conservative benchmark for assessing when their market exclusion is more likely.

Short-lived buyers. Buyers are exposed only to the idiosyncratic risk of buying a lemon. There is no aggregate uncertainty as in Wolinsky (1990) or Blouin and Serrano (2001), in which all units of the good have the same quality and the main trade-off for buyers is between delaying trade to acquire more information or trading early at a potentially larger loss. In my setup, short-lived buyers are not essential to the main model insights and they simplify exposition. Moreover, a long-lived buyer may extract some trade surplus as his bargaining position is likely to strengthen. As for the bargaining protocol, I assume short-lived buyers because it seems a conservative choice to study when high quality sellers are more likely to stay out of the market.

1.4 Equilibrium analysis

Before presenting the main results in sections 1.4.1 and 1.4.2, I restrict attention to economies in which the temporal dimension may help to alleviate the adverse selection problem. Formally, I do *not* consider economies where all sellers can trade immediately, because no allocative efficiency problem arises. It is straightforward to realize that this equilibrium outcome exists only if all sellers post a price p and buyers always accept. In this pooling equilibrium, (i) buyers accept p even when $\xi = L$, and (ii) high quality sellers find it profitable to participate in the market. This equilibrium requires $p \leq \mathbb{E}_{\pi_B}[\theta | \mathscr{I}_B(p, 0, L)]$ and $p \geq v_H + c$. These two conditions imply

[30] For example, it is also the case for $\delta^{\kappa} p - v_{\lambda}$, a utility specification that can be easily interpreted as a seller who discounts future prices but incurs the production cost v_{λ} before entering the market.

$\mathbb{E}_{\pi_B}[\theta|\mathscr{I}_B(p,0,L)] \geq v_H + c$, or:

$$q^0 \geq \frac{\gamma(v_H + c - \theta_L)}{(1-\gamma)(\theta_H - v_H - c) + \gamma(v_H + c - \theta_L)} := q_c^{IP}\text{ [31]}$$

The highest possible pooling price, $p = \mathbb{E}_{\pi_B}[\theta|\mathscr{I}_B(p,0,L)]$, is decreasing in γ; hence q_c^{IP} is increasing in γ. The intuition is straightforward: immediate trade requires buyers to accept when $\xi = L$, and a more informative signal has a stronger negative impact on the posterior expectation. Buyers pay at least $v_H + c$ when they receive $\xi = L$ only if they hold a sufficiently high prior probability q^0 of matching with a high quality seller.

I also refer to another relevant quantity: the minimum value of q^0 above which all sellers trade in a static version of the model with *no* informative signals. As this is simply Akerlof (1970) model with an initial search cost c, it is easy to conclude all sellers trade only if $q^0 \geq \frac{v_H - \theta_L + c}{\theta_H - \theta_L} := q_c^S$. This quantity is a benchmark for understanding how the temporal dimension may improve market outcomes relative to a static model. In the remainder of this paper, unless specified, I assume $q^0 < q_c^S$.

1.4.1 Separating equilibria

A first natural question concerns the existence of a separating equilibrium. The literature shows that it is possible for sellers to separate overtime. In this respect, waiting is a signalling device analogous to education in the classic Spence (1973) model. Buyers find this separating mechanism credible, and they are willing to pay higher prices for sellers who have been on the market longer. Importantly, inter-temporal separation works even when buyers do not have any informative signal ($\gamma = \frac{1}{2}$).

A few common assumptions make ITS possible: (i) sellers' delay costs enter utility through discounting (at rate $\delta < 1$); (ii) there exist strictly positive gains from trade for all types of goods; and (iii) sellers discount an instantaneous payoff as sale and production occur contemporaneously, or goods are durable.[32]

Assumption (i), (ii) and (iii) lead to a preference specification $\delta^\kappa(p - v_\lambda)$ whenever a seller trades after κ previous periods in the market. For a utility function $u(\lambda,\kappa,p)$, the strict single-crossing condition is satisfied if $\frac{u_p(\lambda,\kappa,p)}{|u_\kappa(\lambda,\kappa,p)|}$ is strictly increasing in λ and it has the same sign for all (λ,κ,p) (see Milgrom and Shannon (1994)). Payoff discounting has $u(\lambda,\kappa,p) = \delta^\kappa(p - v_\lambda)$, and the ratio of partial derivatives is $\frac{\delta^\kappa}{|(p-v_\lambda)\delta^\kappa \ln\delta|}$, is always positive and is strictly increasing in λ as $v_H > v_L$. In con-

[31] The subscript c indexes threshold values for q^0 to the search cost c. Later I use q_0^{IP} to denote the value of the threshold for $c = 0$.

[32] A durable good provides a per period flow utility y to its owner, so it is worth $\frac{y}{1-\delta}$ to the seller. Alternatively, sale and production occur contemporaneously when a seller can produce, at the time of trade, a good at cost v_λ.

trast, in my model $u(\lambda,\kappa,p) = p - v_\lambda - (\kappa+1)c$ and this ratio is equal to $\frac{1}{c}$. Similarly, if $u(\lambda,\kappa,p) = \delta^\kappa p - v_\lambda$, then $\frac{\delta^\kappa}{|p\delta^\kappa \ln\delta|}$ is constant for all λ. The last specification describes an economy where sellers pay the production cost before market participation.

Payoff discounting of an instantaneous payoff—as used in models with ITS—makes sellers' individual rationality constraint redundant. As long as they can trade at a price greater or equal to their reservation value v_λ, they can wait indefinitely since $\delta^\kappa(p - v_\lambda) \geq 0$. Indeed, previous models share the feature that H-sellers only trade at $p = v_H$, de facto eliminating their temporal preferences. This property is crucial for ITS; in fact, high quality sellers accommodate any period of delay deemed necessary to prevent low quality sellers from deviating.[33] Final allocations are inefficient because some sellers delay their trades, but all sellers eventually trade.

In my setup the ITS mechanism is not possible without informative signals.

Proposition 1.4.1 *Assume time on market is observable and buyers have no signals. For every $c > 0$ an inter-temporal separating equilibrium does not exist.*

Intuitively, time could credibly signal higher quality only if H-sellers incur a cumulative utility loss larger than all gains from trade for high quality goods. Therefore, market participation is no longer profitable, and they prefer to stay out of the market. Formally, the incentive compatible delay period leads to a utility loss, which violates the individual rationality constraint of H-sellers.[34]

Nonetheless, separation is possible thanks to the combined effect of buyers' informative signals and search costs. I first introduce a notion of separating equilibrium.

Definition 1.4.1 *An equilibrium assessment is separating if H- and L-sellers post different prices after every history $h_i^\kappa \in H^\kappa$ and buyers accept with positive probability.*

Proposition 1.4.2 characterizes the unique separating equilibrium of the game.

Proposition 1.4.2 *Irrespective of time on market observability, a separating equilibrium exists if and only if*

$$c \in \left[\frac{1-\gamma}{\gamma}(\theta_H - \theta_L), \gamma(\theta_H - v_H)\right]$$

In equilibrium, high quality sellers post θ_H and low quality sellers post θ_L. Buyers accept θ_H only after a high signal, but they always accept θ_L.

[33] A hybrid model with search costs c and payoff discounting $\delta < 1$ may or may not admit the existence of an ITS equilibrium. For every c, it is possible to find a δ sufficiently close to 1 such that ITS is not possible; see Example 1.8.1 in Appendix A. I focus on the $\delta = 1$ case because it provides a neat characterization of the results and a sharp intuition of the underlying economic mechanism.

[34] The impossibility of this result depends on the additive specification of delay costs, and it is unchanged even if I consider the set of equilibria in which buyers may play fully mixed strategies (see the proof of Propositon 1.4.1).

A separating equilibrium exists when both search costs and signal precision are sufficiently *high*. If signal precision is too low, the interval in Proposition 1.4.2 does not exist. A low quality seller does not post θ_H, and wait until a buyer receives a high signal, because he expects to pay too high a search cost compared to the immediate payoff of revealing his type and trading at a lower price θ_L. On average, H-sellers receive a high signal after $\frac{1}{\gamma}$ periods, while L-sellers do so after $\frac{1}{1-\gamma}$ periods; as a result, informative private signals create an asymmetric cost of delay between seller' types. Importantly, no separating equilibrium exists when signals are uninformative ($\gamma = \frac{1}{2}$). In other words, differences in the probability of receiving a high signal restore a single-crossing condition and allow separation. The temporal dimension is a necessary condition, but it contributes to the creation of a credible signalling device only with sufficiently informative signals and an adequate level of search costs.[35]

1.4.2 Pooling equilibria

This section provides a complete characterization of equilibria when the search cost is close to zero. I separately analyze each public information setup—TMO and TMN—although the underlying economic intuition is similar.

When c is small, low quality sellers have strong mimicking incentives. Lemma 1.4.1 states the admissible behavioural strategies on the equilibrium path.

Lemma 1.4.1 *There exists $c^* > 0$ such that for every $c \leq c^*$ every equilibrium path only admits the following behavioural strategies for sellers in S^κ:*

- *TMO and $q^0 < q_c^S$:*
  - *H- and L-sellers post a price p^κ that buyers accept only after a high signal.*
  - *H- and L-sellers only post prices rejected with probability one.*[36]
  - *H-sellers stay out of the market and L-sellers trade at price θ_L.*
- *TMN and $q^0 < q_c^{IP}$:*

[35] Proposition 1.4.2 is a pure strategy PBE. Definition 1.3.1 allows mixed strategies for sellers, but it excludes them for buyers. If buyers could play mixed strategies, it is easy to show that a separating equilibrium in mixed strategies would exist for all $c > 0$ if and only if $\gamma \geq \frac{\theta_H - \theta_L + c}{2\theta_H - v_H - \theta_L + c} > \frac{1}{2}$. After receiving a high signal, buyers should play a specific randomization between accepting or rejecting θ_H. A lower search cost requires a higher rejection probability. However, this mixed strategy equilibrium does not have a very realistic flavour. It suggests that in real-world markets buyers play sophisticated randomizations with the sole purpose of helping to separate sellers. Keeping in mind this theoretical possibility, I prefer to focus on equilibria in which buyers play pure strategies and separation is mainly determined by sellers' behaviour, search costs, and the information structure. Lastly, even with mixed strategies, separation is not possible in the limit case of $c = 0$, while the equilibria in section 1.4.2 continue to exist.

[36] This case is not interesting, and it is a pathological result of signalling games. The PBE notion allows for these 'sudden stops' in trade when buyers hold pessimistic beliefs on sellers in cohort S^κ and accept only if $p \leq \theta_L$; in equilibrium, both types of seller prefer not to trade and move to period $\kappa + 1$. These behavioural strategies are ruled out by the undefeated refinement. I do not use the refinement at this stage to stress that the pooling result does not rely on the refinement or other specific restrictions on out-of-equilibrium beliefs.

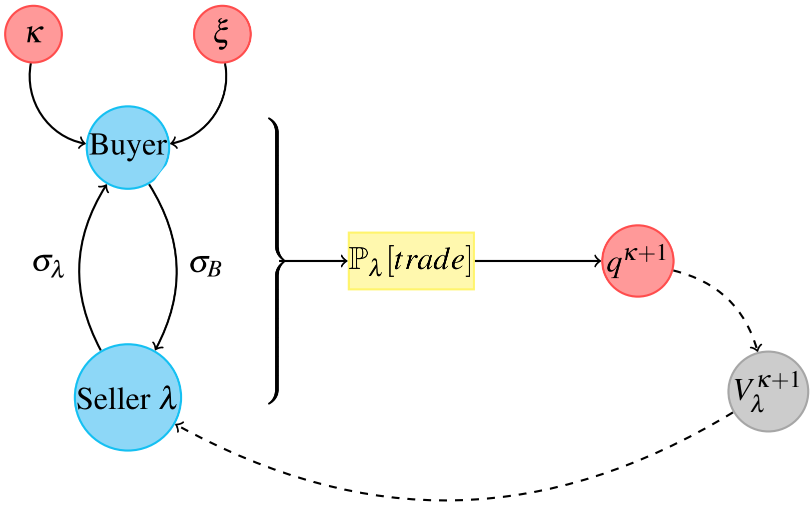

Figure 1.1: Equilibrium interactions with TMO.

- *H- and L-sellers post a price $\bar{p}$ that buyers accept only after a high signal.*
- *H-sellers stay out of the market and L-sellers trade at price θ_L.*

Irrespective of time on market observability, if H-sellers participate in the market—posting a price accepted with positive probability—equilibria are possible only in pooling strategies.[37] Both types of seller offer the same price and buyers accept only if they receive a high signal. When time on market is observable, prices may be different among different cohorts of sellers, but they post a unique price when this information is not available.[38] Even if a low quality seller is unlikely to receive a high signal, this event has a strictly positive probability $1-\gamma > 0$. Low search costs reduce the cost of finding a new buyer, and low quality sellers find it profitable to demand a high price, looking for a buyer who receives a wrong signal. If they happen not to sell at a high price, they can always reveal their type and trade at θ_L. The absence of a credible signalling device precludes separation, and different types of sellers pool on the same action.

Lemma 1.4.1 helps to characterize the set of equilibria of this dynamic signalling game; in the TMO case, the model is non-stationary because equilibrium strategies may depend on κ. Indeed, this dynamic game cannot be solved recursively, as current strategies depend on future continuation values and, viceversa, the latter generally depend on the former because they endogenously determine the share of high quality sellers in each cohort. Figure 1.1 represents this equilibrium interaction between current strategies and future continuation values. To further illustrate this point, consider how buyers form expectations. First, they have a prior probability of being randomly

[37] The *if* clause is crucial. Propositions 1.4.3 and 1.4.4 show that H-sellers may *not* participate in the market.
[38] Lemma 1.4.1 has different thresholds for q^0; see the proof of Lemma 1.4.1 for further details.

matched with a high quality seller in S^κ. In equilibrium, this prior is determined endogenously and is equal to the share of H-sellers in cohort S^κ because of uniform random matching (i.e. $q^\kappa = \frac{S_H^\kappa}{S^\kappa}$). Once matched with seller, a buyer observes his age and posted price, and she updates her beliefs according to signal ξ. The value q^κ plays a substantial role in forming expectations, and it contributes to determine the maximum price that buyers accept from a seller belonging to cohort S^κ. In turn, equilibrium prices determine whether H-sellers want to participate in the market, while equilibrium strategies determine the type-dependent trade probabilities and the evolution of q^κ across different cohorts S^κ.

Due to Lemma 1.4.1, the set of admissible behavioural strategy is tractable—either a pooling price accepted after a high signal or only L-sellers trade—and the equilibrium characterization is straightforward. The next two subsections illustrate how final market outcomes depend on time on market observability.

Time on market observability

Time on market observability (TMO) refers to buyers' ability to observe how long each seller has been participating in the market. Although this information is specific to each individual seller, it plays a crucial role in shaping overall market dynamics. The model provides a tractable framework to analyze how the bilateral asymmetric information problem affects aggregate market dynamics and—reciprocally—how market dynamics influence the possible terms of trade in bilateral transactions.

Proposition 1.4.3 characterizes the unique undefeated equilibrium of the game for c sufficiently small.

Proposition 1.4.3 *Let $q^0 < q_c^S$. There exists $c^* > 0$ such that $\forall c \leq c^*$ there is a unique undefeated equilibrium.*

1. *If* $q^0 \geq q_c^O := \dfrac{(1-\gamma)(v_H + \frac{c}{\gamma} - \theta_L)}{\gamma(\theta_H + \frac{c}{\gamma} - v_H) + (1-\gamma)(v_H - \theta_L - \frac{c}{\gamma})}$

   - *For $\kappa \leq \kappa^*(q^0) < \infty$, $\kappa^*(q^0) = \max\left\{\kappa \in \mathbb{N}_0 : q^\kappa \geq q_c^O\right\}$, H- and L-sellers in S^κ post*
   $$p^\kappa = \mathbb{E}_{\pi_B}[\theta | \mathscr{I}_B(p^\kappa, \kappa, H)]$$
   and buyers accept if and only if $\xi = H$.
   - *After $\kappa^*(q^0) + 1$ periods H-sellers exit the market while L-sellers post θ_L and trade.*

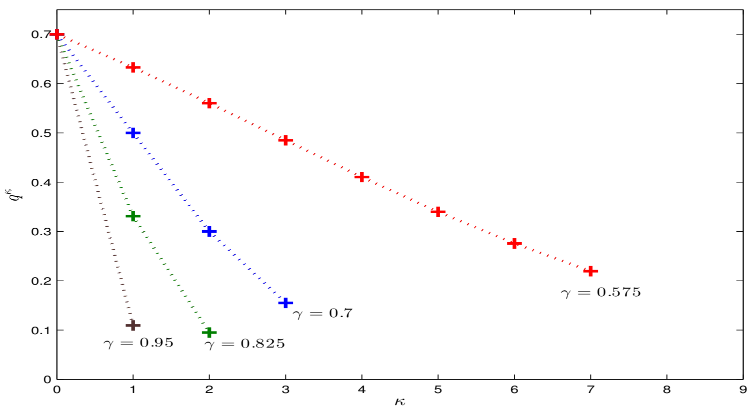

Figure 1.2: Number of periods on the market for H-sellers $\kappa^*(q^0)$ and signal precision γ.

- *For $\kappa < \kappa^*(q^0)$ the share of H-sellers across different cohorts is decreasing in κ:*

$$q^{\kappa+1} = \frac{(1-\gamma)q^\kappa}{(1-\gamma)+(2\gamma-1)(1-q^\kappa)}$$

2. *If $q^0 < q_c^O$ only L-sellers participate in the market and trade at price θ_L.*

Proposition 1.4.3 proves the existence of a unique undefeated equilibrium for sufficiently small c. All sellers from cohort S^κ post the same price p^κ, and buyers accept only if they receive a high signal $\xi = H$. H-sellers are more likely to trade and, on average, they exit the market more rapidly than L-sellers. Thanks to the law of large numbers, q^κ can be expressed as the solution of a first order difference equation. The share of H-sellers is decreasing in κ: the longer a seller has been on the market, the lower buyers' prior belief to match with a high quality seller. Once this belief falls below the minimum threshold q_c^O, no buyer would be willing to pay a price above H-sellers' reservation utility, and the latter prefer to drop out of the market.

Taylor (1999) is the first to point out a negative price externality on older cohorts of sellers. Recently, Kaya and Kim (2014) obtain a similar dynamic when the initial prior belief in meeting a high quality seller is sufficiently high. Proposition 1.4.3 suggests a declining price path *and* a decision to exit the market after a finite number of periods. No market dropout occurs in Taylor (1999) or Kaya and Kim (2014). My result differs from Kaya and Kim's when the prior probability is low ($q^0 < q_c^O$): they predict a dynamic closely related to the ITS mechanism, while Proposition 1.4.3 suggests that H-sellers stay out of the market.

Greater signal precision γ leads to a more rapid decrease in q^κ (see Figure 1.2). However, γ's effect on the measure of H-sellers that exit the market without trading is ambiguous: an increase in γ may reduce $\kappa^*(q^0)$ but it also increases the share γq^κ of

H-sellers who trade in every period $\kappa \leq \kappa^*(q^0)$.

Time on market not observable

Even when time on market is not observable, for small c Lemma 1.4.1 guarantees that all sellers post the same price and buyers accept only if they receive a high signal. Buyers do not distinguish sellers' cohorts, so their prior belief in matching with a high quality seller does not depend on κ and is equal to the share of H-sellers in the overall market. In a stationary equilibrium, this share does not change over time because the mass of each type of seller is constant, i.e. $\bar{q}_t = \bar{q}$ and $S^{\kappa}_{\lambda,t} = S^{\kappa}_{\lambda}$ for every t and λ. This is possible if and only if the entry and exit flows are equal for each type. The entry and exit conditions impose a pair of equations that jointly determine $\overline{q}$ and the overall measure of sellers, say $\overline{S}$.

$$\begin{aligned} &\text{H-sellers:} \quad q^0 = \overline{S}\gamma\overline{q} \\ &\text{L-sellers:} \quad (1-q^0) = \overline{S}(1-\gamma)(1-\overline{q}) \end{aligned}$$

The following proposition describes the equilibrium.

Proposition 1.4.4 *Let $q^0 < q_c^{IP}$. There exists $c^* > 0$ such that $\forall c \leq c^*$ there is a unique undefeated equilibrium.*

1. *If* $q^0 \geq q_c^N := \dfrac{v_H - \theta_L + \frac{c}{\gamma}}{\theta_H - \theta_L} > q_c^S$
   - *Both types of sellers post price*
   $$p^k = \overline{p} = \mathbb{E}_{\pi_B}[\theta | \mathscr{I}_B(\overline{p}, H)] = q^0\theta_H + (1-q^0)\theta_L$$
   for all $\kappa \in \mathbb{N}_0$ and buyers accept if and only if $\xi = H$.
   - *In every period*
   $$\bar{S} = \frac{\gamma - q^0(2\gamma - 1)}{\gamma(1-\gamma)} \qquad \overline{q} = \frac{q^0(1-\gamma)}{\gamma - q^0(2\gamma-1)} < q^0$$
2. *If $q^0 < q_c^N$ only L-sellers participate in the market and they trade at price θ_L.*

Similar to Proposition 1.4.3, H-sellers do not participate in the market when their initial share is too small ($q^0 < q_c^N$). In comparison to the TMO case, they participate in the market for a smaller set of economies as the threshold q_c^N is strictly higher than q_c^O.[39] If the initial share of high quality sellers is sufficiently high, all sellers post

[39] Precisely, this holds when $\frac{c}{\gamma} < \theta_H - v_H$. This is a necessary condition to have $q_c^N < 1$.

a unique price $\overline{p}$ irrespective of their previous periods in the market. Sellers are not penalized if they trade late, because buyers do not observe previous time on market, and they hold a single prior probability $\bar{q}$. If a high quality seller participates in the market, he will continue to do so until he trades. This result follows directly from the forward-looking maximization problem and the fact that previous search costs are sunk.

The equilibrium share of H-sellers $\bar{q}$ is strictly lower than q^0. The underlying economic intuition is simple: on average, H-sellers trade before L-sellers ($\frac{1}{\gamma}$ versus $\frac{1}{1-\gamma}$ periods, respectively); the latter stay longer on the market and decrease H-sellers' market share below q^0. The value of $\bar{q}$ is negatively related to signal precision γ because it decreases the average time on market for H-sellers and increases it for L-sellers. The negative impact of γ on $\bar{q}$ perfectly outweighs the positive effect that a higher signal precision has on buyers' posterior beliefs after $\xi = H$. This feedback effect makes signal precision irrelevant for equilibrium prices; in fact, the pooling price $\overline{p} = q^0\theta_H + (1-q^0)\theta_L$ does not depend on γ. In particular, it is equal to buyers' expected value for a good offered by newly born sellers (S^0) *before* receiving a signal.

When $\gamma > \frac{1}{2}$ all sellers trade immediately only if $q^0 \geq q_c^{IP} > q_c^S$. Proposition 1.4.4 implies that H-sellers participate only if $q^0 \geq q_c^N > q_c^S$. Therefore, for small c neither the temporal dimension nor buyers' informative signals mitigate the adverse selection problem. Actually, because $q_c^N > q_c^S$, all sellers trade for a strictly smaller set of economies compared to the classic static adverse selection model (although $q_0^N = q_0^S$). Janssen and Roy (2004) point out that the infinite repetition of the static equilibrium is the only stationary equilibrium of a dynamic adverse selection model with uninformed buyers. Proposition 1.4.4 suggests that this result also applies when buyers have informative signals. Even though this conclusion seems extreme, the mechanism in place is suggestive and it might be worth assessing its empirical validity.

1.5 Welfare analysis

In this section, I first study how different information structures compare in terms of welfare. Then, I discuss how a small c may actually reduce H-sellers' market participation compared to a situation with higher costs.

Definition 1.5.1 introduces a simple notion of allocative efficiency that does *not* take into account how many periods elapse before sellers trade.

Definition 1.5.1 *An equilibrium is allocative efficient if all sellers trade in finite time almost surely.*

In my model setup with transferable utility, it is natural to adopt a utilitarian welfare criterion. Clearly, if an equilibrium is not allocative efficient then it does not maximize total welfare. As there is no time discounting, if $c \searrow 0$ an allocative efficient equilibrium is arbitrarily close to maximize utilitarian welfare. Moreover, since the transfer scheme is budget balanced—prices paid by buyers are equal to prices received by sellers—a utilitarian allocation is also Pareto efficient.

Proposition 1.5.1 summarizes the welfare properties of equilibria in the relevant set of economies with $q^0 < q^S_c$.

Proposition 1.5.1 *Let $q^0 < q^S_c$. The following statements hold when $c \searrow 0$:*

1. *No equilibrium allocation maximizes total welfare.*
2. *If $q^0 \in \left[q^O_0, q^S_0\right)$ TMO improves welfare compared to TMN.*
3. *Under TMN welfare is identical to the static Akerlof (1970) model.*

When time on market is not observable and $q^0 < q^N_0 = q^S_0$, only L-sellers trade and final allocations are identical to the ones in Akerlof (1970) model. Obviously, the equilibrium allocation does not maximize total welfare. When time on market is observable and $q^0 \in [q^O_0, q^S_0)$, all H-sellers initially participate in the market, but a strictly positive measure does not trade their goods because they drop out after a finite number of periods (see Proposition 1.4.3). Not all mutually beneficial exchanges take place, resulting in allocative inefficiency. Nevertheless, H-sellers participate and trade with positive probability at least for one period, but they always stay out of the market if time on market is not observable. To sum up, for small c, a dynamic model with private informative signals achieves a welfare improvement compared to a static model with uninformed buyers only when time on market is observable and $q^0 \in [q^O_0, q^S_0)$.

Figure 1.3 illustrates the results of Proposition 1.5.1. Each economy is parametrized by a (γ, q^0) coordinate. Depending on time on market observability, different equilibria exist: immediate pooling (white; $q^0 \geq q^{IP}_0$), pooling with TMN (white and blue; $q^0 \geq q^S_0$), pooling with TMO[40] (white, blue and green; $q^0 \geq q^O_0$), and exclusive market participation by L-sellers irrespective of time on market observability (red; $q^0 < q^O_0$).

Predicting whether H-sellers are more likely to participate in the market when c is "small" or "large" is not clear cut. Naively, lower search costs increase final payoffs and relax their individual rationality constraint. However, this intuition does not account for how equilibria might change.

On the one hand, if only pooling equilibria existed, a smaller c would enlarge the set of economies in which H-sellers participate in the market; in fact, both q^O_c and q^N_c are

[40] For $c < c^*$ the equilibria in Proposition 1.4.3 and 1.4.4 exist for every $q^0 \geq q^O_0$ and $q^0 \geq q^N_0$, respectively. However, they may not be unique when $q^0 \geq q^S_0$ or $q^0 \geq q^{IP}_0$, respectively.

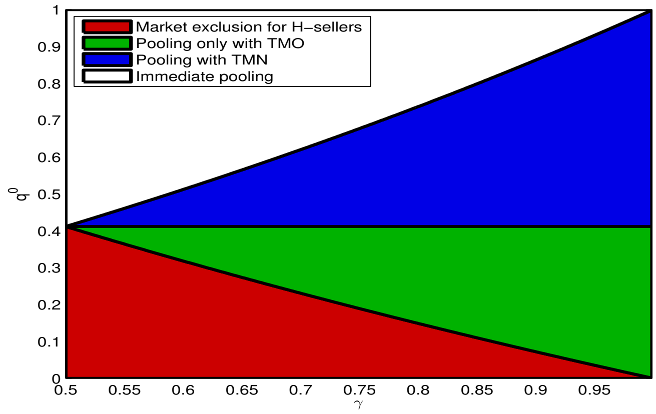

Figure 1.3: Equilibria in different economies $\mathscr{E}(\gamma, q^0)$ for $c \leq c^*$.

decreasing in c. On the other hand, a small c encourages L-sellers to post high prices and take advantage of buyers' imperfect signals. If q^0 is below q_c^O or q_c^N H-sellers stay out of the market. Furthermore, pooling prices convey no additional information on the underlying assets, so they rank at the bottom in terms of informational efficiency.

Proposition 1.4.2 characterizes the separating equilibrium and it suggests that a small c may reduce trade opportunities when signal precision is sufficiently high. Indeed, for every q^0 a separating equilibrium exists only if:

$$c \in \left[\frac{1-\gamma}{\gamma}(\theta_H - \theta_L), \gamma(\theta_H - v_H)\right]$$

In a separating equilibrium all sellers participate in the market and trade. Final allocations do not maximize welfare because sellers pay strictly positive search costs, but all sellers eventually trade.

The beneficial signalling effect of search costs may extend to "intermediate" values of c, i.e. when search costs are too small to create a separating equilibrium but too large to support a pooling equilibrium. Unfortunately, providing a complete equilibrium characterization for all values of c is a complex endeavor, especially in the TMO case. As a result, I justify this claim through a specific semi-separating equilibrium.

Proposition 1.5.2 *Irrespective of time on market observability, there exists a region of parameters* $(\theta_H, v_H, \theta_L, v_L, q^0, \gamma)$ *where*

- *Only L-sellers trade for sufficiently small c.*
- *For*
$$c \in \left[\frac{\gamma(1-\gamma)}{\gamma^2+\gamma-1}(v_H - \theta_L), \frac{1-\gamma}{\gamma}(\theta_H - \theta_L)\right)$$
there exists at least one semi-separating equilibrium in which all sellers trade.

In this semi-separating equilibrium all H-sellers participate in the market and trade. However, γ and c should be sufficiently high to exclude complete pooling on the same action. In equilibrium, H-sellers only post a high price and L-sellers mix between the high price and θ_L. Posted prices do not depend on κ and all sellers trade over time.[41]

To sum up, search costs can be beneficial by discouraging low quality sellers from pretending to have a high quality good. Although a small c makes participation cheaper, it may worsen adverse selection, and leave high quality goods out of the market when their initial share is low. In the next section, I consider whether it is possible to enjoy the welfare benefits of low search costs without exacerbating the adverse selection problem.

1.6 Market design

As previously explained, when the cost c is small H-sellers may have less incentive to participate in the market. Even when all sellers participate, equilibria are in pooling strategies and prices do not provide any information on product quality. If informational efficiency is considered relevant, this is another loss to take into account.

From a policy perspective, understanding how a benevolent market designer can intervene to promote full market participation for the largest possible set of economies is crucial. I adopt a stringent benchmark for the objectives of the market design intervention: the resulting equilibrium has to achieve both allocative (see Definition 1.5.1) and informational efficiency (i.e. prices reveal sellers' types). In my setup, an allocative efficient equilibrium maximizes utilitarian welfare when $c = 0$.

The market designer is subject to a series of reasonable limitations. First, the mechanism has to be budget balanced on the equilibrium path. This restriction seems natural as the market should not depend on any external amount of resources to induce participation and trade. Second, transfers cannot be conditional on any posted price. Differently from buyers, the market designer cannot observe currently posted prices. This restriction is consistent with the idea that bilateral transactions involve elements of private negotiation that are difficult to verify externally.[42] Therefore, transfers can only be conditional on market participation (τ), trade (r) and, possibly, time on market (κ). If time on market is observable, transfers (τ^κ, r^κ) can vary across different sellers' cohorts. If instead time on market is not observable, transfers are constant, i.e. $(\tau^\kappa, r^\kappa) = (\tau, r)$. I consider a budget balanced mechanism that satisfies these

[41] See the proof of Proposition 1.5.2 for a complete characterization of the equilibrium.

[42] For example, parties may exchange side payments in order to misreport posted prices. Setting up a market mechanism with price contingent transfers and robust to side payments goes beyond the scope of this paper. Moreover, if a designer could observe currently posted prices, he could reconstruct the price history for each seller and would have more information than buyers.

properties to be *feasible*. For every cohort of sellers, a feasible mechanism has to be ex interim individually rational as a seller knows his type when he participates in the market. I consider a feasible mechanism to be *efficient* if it leads to an allocative and informationally efficient equilibrium.

Proposition 1.6.1 *Let $c = 0$. When time on market is observable an efficient mechanism exists only in economies with*

$$q^0 \geq q^* := \max\left\{0, 1 - \left(\frac{\gamma}{1-\gamma}\right)^2 \frac{\theta_H - v_H}{\theta_H - \theta_L}\right\}$$

The efficient market mechanism implements a separating equilibrium with:

- *a constant market participation tax $\tau^* = \frac{1-\gamma}{\gamma}(\theta_H - \theta_L)$.*
- *a fixed tax rebate $r^* = \tau^*\left(1 + \frac{1-\gamma}{\gamma}q^*\right)$ once the seller trades.*

For $q^0 < q^*$ no feasible mechanism improves market outcomes and only L-sellers trade.

In equilibrium, a low quality seller posts θ_L and buyers accept this price for every signal realization, while a high quality seller posts θ_H and trades once he matches with a buyer who receives a high signal. Prices reveal sellers' types and the equilibrium is informationally efficient.[43] The efficient market intervention is invariant with respect to κ because transfers do not depend on cohort S^κ. This is related to the fact that prices reveal types, and information on the specific cohort becomes irrelevant for inferring product quality. As (τ^*, r^*) is κ-invariant, the same mechanism is efficient when time on market is not observable.

The green and yellow areas in Figure 1.4 illustrate the improvement due to (τ^*, r^*). Without a market intervention, an equilibrium is allocative efficient only if $q^0 \geq q_0^S$ (blue and white areas). When time on market is observable, H-sellers could also participate in the market in economies in the green area $q^0 \in [q_0^O, q_0^S)$, but they would not trade for sure. No high quality seller participates in economies in the yellow and red areas.

Proposition 1.6.1 points out that the mechanism (τ^*, r^*) may support an allocative and informational efficient allocation for every$q^0 \in (0,1)$ only if:

$$\gamma \geq \frac{\sqrt{\theta_H - \theta_L}}{\sqrt{\theta_H - \theta_L} + \sqrt{\theta_H - v_H}} := \gamma^*$$

[43] In supplementary work, I relax the requirement of informationally efficient prices and explore whether it is possible to implement an allocative efficient equilibrium for an even larger set of economies. I find that no improvement is possible when time on market is not observable. However, I could not prove an analogous general result when time on market is observable; nonetheless, I could not find any counterexample.

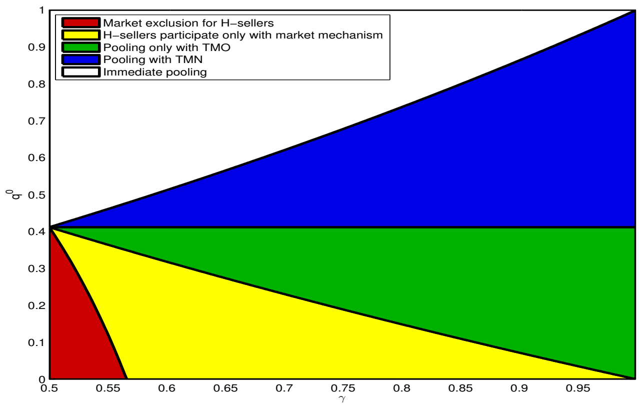

Figure 1.4: Efficient market intervention (τ^*, r^*).

Despite the improvement, it is still not possible to implement a first best allocation in every economy. If $\gamma < \gamma^*$ an efficient allocation is possible only if $q^0 \geq q^* > 0$ (see the red area in Figure 1.4).[44] Otherwise, it is not possible to mitigate adverse selection with a feasible market intervention. In these economies, the mechanism (τ^*, r^*) violates the individual rationality constraint of H-sellers because of the budget balance restriction. High and low quality sellers have different expected time on market ($\frac{1}{\gamma}$ and 1 periods respectively), and budget balance leads to an implicit transfer from high to low quality sellers, reducing the former's expected payoff. A high quality seller prefers to stay out of the market when this expected transfer outweighs his gains from trade.

1.7 Conclusion

Several theoretical papers point out the existence of an ITS equilibrium in markets with asymmetric information on asset quality. All sellers participate in the market: low quality sellers trade early, and high quality sellers trade later at higher prices. Since most markets offer multiple opportunities to sell, these results seem favorable compared to Akerlof's conclusions. However, a few recent empirical contributions document price patterns that conflict with the implications of ITS.

I study a dynamic adverse selection model with search costs: sellers incur symmetric and additive costs when searching for a new trade opportunity, and the market participation decision is non-trivial. The resulting analysis provides two main benefits. First, it highlights the role of search costs in the provision of a credible signalling device; second, it suggests market design policies to enhance participation in markets—such as online trading platforms—that currently look for ways to attract high quality

[44] Notice that $\gamma^* < \frac{\theta_H - \theta_L}{2\theta_H - v_H - \theta_L}$; see footnote 35 for a definition. As a consequence, the mechanism also improves on a hypotetical separating equilibrium in which buyers are allowed to use mixed strategies.

products.

I present a framework with only two types of goods and binary signals. Despite its simplicity, the model allows me to uncover the main economic mechanisms at work. Future extensions may broaden the setup to multiple goods and more general signal distributions. I expect that the main economic intuition would continue to hold. Another extension would be to introduce heterogenous buyers who search for different product qualities in a directed search environment.[45]

1.8 Appendix A

1.8.1 Extended notation

A behavioural strategy for seller $i \in S_\lambda^\kappa$ is a function $\sigma_{\lambda,i}^\kappa : H^{\kappa-1} \to \Delta(A_S)$. Let $\Sigma_{\lambda,i}^\kappa$ be the set of all strategies $\sigma_{\lambda,i}^\kappa$ (let $\Sigma_{\lambda,i} = \bigcup_{\kappa=0}^{\infty} \Sigma_{\lambda,i}^\kappa$ and $\sigma_{\lambda,i} \in \Sigma_{\lambda,i}$). A behavioural strategy for a buyer matched with seller i is a function $\sigma_{B,i} : \mathscr{I}_B \to \Delta(A_B)$, where $\mathscr{I}_B$ denotes her information set. It is $\mathscr{I}_B(p,\kappa,\xi)$ if time-on market is observable and $\mathscr{I}_B(p,\xi)$ if it is not. When it is not relevant to specify whether TMO or TMN applies I denote an information set with $\mathscr{I}_B(p,\cdot,\xi)$.

Let $\pi_{\lambda,i}(\xi|h_i^{\kappa-1})$ be seller's $i \in S_\lambda^\kappa$ belief that his matched buyer receives signal ξ after history $h_i^{\kappa-1}$. I denote with $\mathbb{P}^{\sigma^*\pi^*}(z|h_i^{\kappa-1})$ the probability to reach terminal history $z \in Z_i$ from $h_i^{\kappa-1} \in H^{\kappa-1}$ under the assessment (σ^*,π^*). I use $V_{\lambda,i}^\kappa(\sigma^*,\pi^*|h_i^{\kappa-1})$ for the continuation value to a seller $i \in S_\lambda^\kappa$ with previous history $h_i^{\kappa-1}$. It uses the utility function $u_\lambda^\kappa(z) := \tilde{u}_\lambda(z) + \kappa c$, which ignores the previous κc sunk costs.

Each seller maximizes his inter-temporal expected utility after every history $h_i^{\kappa-1}$, $\kappa \in \mathbb{N}_0$.

$$V_{\lambda,i}^\kappa(\sigma_i^*,\sigma_{-i}^*,\pi_i^*,\pi_{-i}^*|h_i^{\kappa-1}) = \max_{\sigma_i \in \Sigma_{\lambda,i}} \sum_{z \in \{\cup Z_i^j\}_{j=\kappa}^{\infty}} \mathbb{P}^{\sigma_i\sigma_{-i}^*\pi_i^*\pi_{-i}^*}(z|h_i^{\kappa-1})u_\lambda^\kappa(z)$$

In equilibrium, if a seller $i \in S_\lambda^\kappa$ posts price p, it is possible to write the value function as:[46]

$$V_{\lambda,i}^\kappa(\sigma^*,\pi^*|h_i^{\kappa-1}) = \sum_{\xi \in \{H,L\}} \pi_{\lambda,i}^*(\xi|h_i^{\kappa-1}) \left[\sigma_B^*(A|\mathscr{I}_B(p,\cdot,\xi))(p-v_\lambda) + \sigma_B^*(R|\mathscr{I}_B(p,\cdot,\xi))V_\lambda^{\kappa+1}(\sigma^*,\pi^*|h_i^\kappa)\right] - c$$

[45] Guerrieri and Shimer (2014) characterize a competitive search equilibrium where high quality sellers separate because they are more willing to accept a lower probability to trade. Jullien and Mariotti (2006) present a similar mechanism for auctions and separation results from setting different type-dependent reservation prices.

[46] Lemma 1.8.1 adapts the results in Hendon et al. (1996) to ensure that the one-shot deviation property holds in this model setup.

A buyer can accept or reject an offer. Her expected payoff is simply:

$$\begin{cases} \mathbb{E}_{\pi_{\mathbb{B}}}[\theta] - p & if\ a_B = A \\ 0 & if\ a_B = R \end{cases}$$

where $\mathbb{E}_{\pi_B}[\theta] = \pi_B\theta_H + (1-\pi_B)\theta_L$ is the expectation under her posterior belief π_B.

Definition 3.3.1 is a formal statement of the equilibrium concept in Definition 1.3.1.

Definition 1.8.1 *An equilibrium of the game with TMO is a stationary and symmetric assessment (σ^*, π^*) such that for every $i \in S_\lambda^\kappa$, $\lambda \in \{H, L\}$, and $\kappa \in \mathbb{N}_0$:*

1. $\sigma_\lambda^{*^\kappa}(a_{S,i}|h_i^{\kappa-1}) \in \arg\max_{\sigma_i \in \Sigma_{\lambda,i}} V_\lambda^\kappa(\sigma_i, \sigma_{-i}^*, \pi^*|h_i^{\kappa-1})\ \forall h_i^{\kappa-1} \in H^{\kappa-1}$.
2. $\sigma_B^*(a_{B,i}|\mathscr{I}_B(p^\kappa, \kappa, \xi))$ *is a pure-strategy best response.*
3. $\pi_B(p, \kappa, \xi)$ *is updated according to Bayes' rule whenever possible.*
4. $\pi_\lambda(\xi|h_i^{\kappa-1}) = \mathbb{P}_\lambda(\xi)$ *for every* $h_i \in H^{\kappa-1}$, $\kappa \in \mathbb{N}_0$ *and* $\lambda \in \{H, L\}$.

The definition slightly restricts the weak Perfect Bayesian Equilibrium concept. Condition 1 allows best response strategies to depend on λ, κ and $h_i^{\kappa-1}$. Condition 2 only considers pure strategy best responses for buyers. Condition 3 requires buyers to update beliefs according to Bayes' rule on the equilibrium path. Since buyers are short-lived, it is not necessary to impose any additional restriction on their out-of-equilibrium beliefs in order to have a *reasonable assessment* (see Definition 3 of Fudenberg and Tirole (1991)). Finally, condition 4 restricts sellers not to change their beliefs on the likelihood that future matched buyers receive signal $\xi \in \{H, L\}$. This restriction seems natural as buyers' signal realizations are independent from H-sellers' previous history. As a result, it is equivalent to a "no signalling what you don't know" condition on sellers' posterior beliefs. Adapting this definition to the TMN setup is straightforward and I omit it in the interest of space. The only difference relates to buyers' impossibility to condition on κ, so they form beliefs using a single prior probability q. I always keep the possibility that behavioural strategy profiles may differ among sellers' cohorts and histories.

Definition 1.8.2 states in the context of my framework the concept of undefeated equilibrium originally presented in Mailath et al. (1993).

Definition 1.8.2 *An equilibrium assessment (σ^*, π^*) defeats $(\tilde{\sigma}, \tilde{\pi})$ if $\exists \{p^\kappa\}_{\kappa=0}^\infty$ such that:*

1. $\exists \kappa \in \mathbb{N}_0$ *with* $\tilde{\sigma}_\lambda^\kappa(p^\kappa|h_i^{\kappa-1}) = 0$ *for every* $\lambda \in \{H, L\}$ *and* $h_i^{\kappa-1} \in H^{\kappa-1}$ *while* $\Lambda(p^\kappa) := \{\lambda : \exists h_i^{\kappa-1} \in H^{\kappa-1}\ s.t\ \sigma_\lambda^{*^\kappa}(p^\kappa|h_i^{\kappa-1}) > 0\} \neq \varnothing$.

2. *$\forall\lambda \in \Lambda(p^\kappa)$ it holds $V^j_\lambda(\sigma^*,\pi^*|h_i^{j-1}) \geq V^j_\lambda(\tilde{\sigma},\tilde{\pi}|h_i^{j-1})$ $\forall h_i^{j-1} \in H^{j-1}$and $j \in \mathbb{N}_0$.*
   Moreover, $\exists\lambda \in \Lambda(p^\kappa)$ s.t. $V^\kappa_\lambda(\sigma^*,\pi^*|h_i^{\kappa-1}) > V^\kappa_\lambda(\tilde{\sigma},\tilde{\pi}|h_i^{\kappa-1})$ for some $h_i^{\kappa-1} \in H^{\kappa-1}$.

3. *$\tilde{\pi}(p^\kappa, q^\kappa, \xi) \neq \frac{\mathbb{P}_H(\xi) q^\kappa \sigma^\kappa_H(p^\kappa)}{\mathbb{P}_H(\xi) q^\kappa \sigma^\kappa_H(p^\kappa) + \mathbb{P}_L(\xi)(1-q^\kappa)\sigma^\kappa_L(p^\kappa)}$ with $\sigma^\kappa_\lambda(p^\kappa)$ satisfying:*

   - *$\lambda \in \Lambda(p^\kappa)$ and $V^\kappa_\lambda(\sigma^*,\pi^*|h_i^{\kappa-1}) > V^\kappa_\lambda(\tilde{\sigma},\tilde{\pi}|h_i^{\kappa-1}) \Rightarrow \sigma^\kappa_\lambda(p^\kappa) = 1$.*
   - *$\lambda \notin \Lambda(p^\kappa) \Rightarrow \sigma^\kappa_\lambda(p^\kappa) = 0$.*

An equilibrium assessment (σ^*,π^*) is undefeated if there is no other equilibrium that defeats (σ^*,π^*) according to Definition 1.8.2.

For notational simplicity I omit to explicitly specify π^* when it is obvious from the context. For instance, I use sometimes $V^\kappa_\lambda(\sigma^*)$ to denote $V^\kappa_\lambda(\sigma^*,\pi^*)$.

1.8.2 Preliminary results

Example 1.8.1 *Inter-temporal separating equilibrium.*

The example explains how ITS works. It is extremely simple and its goal is to make as transparent as possible the main backbone mechanism.

Suppose sellers discount future payoffs at rate δ and they pay a per period search cost c. In equilibrium low quality sellers trade immediately while high quality sellers wait until a future period $t > 0$ to trade at a higher price.[47] This equilibrium exists only if:

$$IR_H: \quad \delta^t(\theta_H - v_H) - \tfrac{1-\delta^{t+1}}{1-\delta}c \geq 0 \qquad \rightarrow \quad t \leq f(\delta,c)$$

$$IC_L: \quad \theta_L - v_L - c \geq \delta^t(\theta_H - v_L) - \tfrac{1-\delta^{t+1}}{1-\delta}c \qquad \rightarrow \quad t \geq h(\delta,c)$$

The functions $f(\delta,c)$ and $h(\delta,c)$ are continuous in both arguments with limits:

$$\lim_{c\searrow 0} f(\delta,c) = +\infty \qquad \lim_{\delta\nearrow 1} f(\delta,c) = \tfrac{\theta_H - v_H - c}{c}$$
$$\lim_{c\searrow 0} h(\delta,c) = \tfrac{1}{\ln\delta}\ln\tfrac{\theta_L - v_L}{\theta_H - v_L} \qquad \lim_{\delta\nearrow 1} h(\delta,c) = \tfrac{\theta_H - v_H}{c}$$

The ITS equilibrium exists only if $f(\delta,c) \geq h(\delta,c)$. This is always the case if $c = 0$ and $\delta < 1$ while it is never so if $c > 0$ and $\delta = 1$ as $v_H > \theta_L$.

[47] To simplify derivation, I assume sellers have full bargaining power but notice that the argument generalizes to other bargaining protocols.

Lemma 1.8.1 *A strategy profile σ_i^* is a sequential best reply to (σ_{-i}^*, π^*) for seller $i \in S$ if and only if $\sigma_i^{*^\kappa}(a_{S,i}|h_i^{\kappa-1})$ is a local best reply to (σ_{-i}^*, π^*) for all $\kappa \in \mathbb{N}_0$ and $h_i^{\kappa-1} \in H^{\kappa-1}$.*

Proof Lemma 1.8.1.

Necessity. It follows directly from the definition of sequential best reply.

Sufficiency. Suppose on the contrary that σ_i^* is a local best reply for every $h_i^{\kappa-1} \in H^{\kappa-1}$ and $\kappa \in \mathbb{N}_0$, but there exists a strategy σ_i' that strictly improves on σ_i^* after history $h_i^{\kappa-1}$. Let this increment be equal to $\varepsilon > 0$. Seller's i expected payoff at $h_i^{\kappa-1}$ is:

$$\begin{aligned} V_{\lambda,i}^{\kappa}(\sigma_i', \sigma_{-i}^*, \pi^*|h_i^{\kappa-1}) \quad &= \sum_{z\in\{\cup Z_i^j\}_{j=\kappa}^{\infty}} \mathbb{P}^{\sigma_i'\sigma_{-i}^*\pi^*}(z|h_i^{\kappa-1})u_\lambda^\kappa(z) = \sum_{z\in Z_i^\kappa} \mathbb{P}^{\sigma_i'\sigma_{-i}^*\pi^*}(z|h_i^{\kappa-1})u_\lambda^\kappa(z) \\ &+ \sum_{h^\kappa\in H^\kappa/Z_i^\kappa} \mathbb{P}^{\sigma_i'\sigma_{-i}^*\pi^*}(h_i^\kappa|h_i^{\kappa-1}) \sum_{z\in\{\cup Z_i^j\}_{j=\kappa+1}^{\infty}} \mathbb{P}^{\sigma_i'\sigma_{-i}^*\pi^*}(z|h_i^\kappa)u_\lambda^\kappa(z) \end{aligned}$$

Let $c(z|h_i^{\kappa-1})$ be the search costs from $h_i^{\kappa-1} \in H^{\kappa-1}$ to terminal history $z \in Z_i^j$, i.e. $c(z|h_i^{\kappa-1}) = (j-\kappa+1)c$.
An upper bound on $\sum_{z\in\{\cup Z_i^j\}_{j=\kappa+1}^{\infty}} \mathbb{P}^{\sigma_i'\sigma_{-i}^*\pi^*}(z|h_i^\kappa)u_\lambda^\kappa(z)$ is:

$$\sum_{z\in\{\cup Z_i^j(p)\}_{j\geq\kappa+1}} \mathbb{P}^{\sigma_i'\sigma_{-i}^*\pi^*}(z|h_i^\kappa)[\theta_H - v_\lambda - c(z|h_i^\kappa)] + \sum_{z\in\{\cup Z_i^j(D)\}_{j\geq\kappa+1}} \mathbb{P}^{\sigma_i'\sigma_{-i}^*\pi^*}(z|h_i^\kappa)[-c(z|h_i^\kappa)] \tag{1.1}$$

Observe that σ_i^* is a local best reply at $h_i^{\kappa-1}$, and a seller can always get a zero continuation value if he drops out of the market. Therefore, for some $h_i^\kappa \in H^\kappa/Z_i^\kappa$ it must be $\mathbb{P}^{\sigma_i'\sigma_{-i}^*\pi^*}(h_i^\kappa|h_i^{\kappa-1}) > 0$ and σ_i' is a profitable deviation only if:

$$\sum_{z\in\{\cup Z_i^j\}_{j\geq\kappa+1}} \mathbb{P}^{\sigma_i'\sigma_{-i}^*\pi^*}(z|h_i^\kappa)u_\lambda^\kappa(z) > \sum_{z\in\{\cup Z_i^j\}_{j\geq\kappa+1}} \mathbb{P}^{\sigma_i^*\sigma_{-i}^*}(z|h_i^\kappa\pi^*)u_\lambda^\kappa(z) \geq -c \tag{1.2}$$

As a result, $\sum_{z\in\{\cup Z_i^j(p)\}_{j\geq\kappa+1}} \mathbb{P}^{\sigma_i'\sigma_{-i}^*}(z|h_i^\kappa) > 0$ otherwise sellers would have the same expected payoff at $h_i^{\kappa-1}$ because σ_i^* is a local best reply. Hence, equations (1.1) and (1.2) imply:

$$(\theta_H - v_\lambda) > \frac{\sum_{z\in\{\cup Z_i^j\}_{j\geq\kappa+1}} \mathbb{P}^{\sigma_i'\sigma_{-i}^*}(z|h_i^\kappa)c(z|h_i^\kappa) - c}{\sum_{z\in\{\cup Z_i^j(p)\}_{j\geq\kappa+1}} \mathbb{P}^{\sigma_i'\sigma_{-i}^*}(z|h_i^\kappa)} \geq \frac{\mathbb{P}^{\sigma_i'\sigma_{-i}^*}(z|h_i^\kappa)c(z|h_i^\kappa) - c}{\sum_{z\in\{\cup Z_i^j(p)\}_{j\geq\kappa+1}} \mathbb{P}^{\sigma_i'\sigma_{-i}^*}(z|h_i^\kappa)}$$

for all $z \in Z_i^j$ and $j \geq \kappa+1$.

Since $c(z \in Z_i^j | h_i^\kappa) \to +\infty$ as $j \to +\infty$ the inequality holds only if $\lim_{j \to +\infty} \mathbb{P}^{\sigma_i' \sigma_{-i}^*}(z \in Z_i^j | h_i^\kappa) = 0$. Therefore, there exists a finite $\hat{t}$ and history $h_i^{\kappa+\hat{t}}$ such that the strategy:

$$\hat{\sigma}_i = \begin{cases} \sigma_i' & \forall j < \kappa + \hat{t} \text{ and } \forall h_i^j \in H^j \\ \sigma_i^* & \forall j \geq \kappa + \hat{t} \text{ and } \forall h_i^j \in H^j \end{cases}$$

improves of at least $\frac{\varepsilon}{2}$ on σ_i^* with a finite number of deviations; i.e:

$$V_{\lambda,i}^\kappa(\hat{\sigma}_i, \sigma_{-i}^* | h_i^{\kappa-1}) - V_{\lambda,i}^\kappa(\sigma_i^*, \sigma_{-i}^* | h_i^{\kappa-1}) \geq \frac{\varepsilon}{2}$$

However, the main result in Hendon, Jacobsen and Sloth (1996) ensures that no finite sequence of deviations can improve on σ_i^*, contradiction. ∎

Proof Proposition 1.4.1.

Suppose per contra there exists an ITS equilibrium (σ^*, π^*). Denote by $\mathbb{P}_\lambda^{\sigma^* \pi^*}(z^j(p) | h_i^{\kappa-1})$ the probability that seller $i \in S_\lambda^\kappa$ reaches the terminal history $z \in Z_i^j(p)$ under this separating equilibrium (σ^*, π^*). For a L-seller $i \in S_L^\kappa$ a deviation strategy $\sigma_i'^j(p | h_i^{j-1}) = \sigma_H^{*j}(p | h_i^{j-1})$, $j \geq \kappa$, for all $h_i^{j-1} \in H^{j-1}$ implies:

$$\mathbb{P}_H^{\sigma^* \pi^*}(z^j(p) | h_i^{\kappa-1}) = \mathbb{P}_L^{\sigma_i' \sigma_{-i}^* \pi^*}(z^j(p) | h_i^{\kappa-1}) \tag{1.3}$$

$\forall z^j(p) \in Z^j(p)$, $j \geq 0$, since signals are not informative ($\gamma = \frac{1}{2}$) and both sellers have identical chances to trade if they play the same strategy profile.

In equilibrium L-sellers do not find strictly profitable to play σ_i' only if:

$$\theta_L - v_L - c \geq \sum_{j=\kappa}^{\infty} \int_0^\infty \mathbb{P}_L^{\sigma_i' \sigma_{-i}^* \pi^*}(z^j(p) | h_i^{\kappa-1}) \, [p - v_L - (j - \kappa + 1)c] \, \mathrm{d}p \tag{1.4}$$

All H-sellers eventually trade, hence:

$$\sum_{j=\kappa}^{\infty} \int_0^\infty \mathbb{P}_H^{\sigma^* \pi^*}(z^j(p) | h_i^{\kappa-1}) \, \mathrm{d}p = \sum_{j=\kappa}^{\infty} \int_0^\infty \mathbb{P}_L^{\sigma_i' \sigma_{-i}^* \pi^*}(z^j(p) | h_i^{\kappa-1}) \, \mathrm{d}p = 1 \tag{1.5}$$

where the first equality follows by equation (1.3). Therefore, L-sellers no deviation condition can be rewritten as:

$$\theta_L - c \geq \sum_{j=\kappa}^{\infty} \int_0^\infty \mathbb{P}_H^{\sigma^* \pi^*}(z^j(p) | h_i^{\kappa-1}) \, [p - (j - \kappa + 1)c] \, \mathrm{d}p \tag{1.6}$$

H-sellers' market participation constraint is satisfied only if:

$$\sum_{j=\kappa}^{\infty} \int_{0}^{\infty} \mathbb{P}_H^{\sigma^*\pi^*}(z^j(p)|h_i^{\kappa-1})\,[p - v_H - (j-\kappa+1)c]\;\mathrm{d}p \geq 0$$

However, by equations (1.5) and (1.6) an upper bound on H-sellers' equilibrium expected payoff is $\theta_L - v_H - c < 0$, a violation of H-sellers' individual rationality constraint. ■

Lemma 1.8.2 *In every equilibrium assessment* (σ^*, π^*) *it holds:*

$$V_\lambda^\kappa(\sigma^*, \pi^*|h_i^{\kappa-1}) = V_\lambda^\kappa(\sigma^*, \pi^*|\tilde{h}_i^{\kappa-1}) := V_\lambda^\kappa(\sigma^*, \pi^*)$$

for every $h_i^{\kappa-1}, \tilde{h}_i^{\kappa-1} \in H^{\kappa-1}$.

Proof Lemma 1.8.2.
Previous histories $h_i^{\kappa-1} \in H^{\kappa-1}$ are not observable to buyers so their best responses with respect to any price offer p from sellers in S^κ are identical for sellers with different histories. Moreover, final payoffs depend only on how many periods κ were previously spent on the market and—in case of trade—on the last price offer p. Therefore, in equilibrium all sellers in S_λ^κ must have the same expected payoff irrespective of $h^{\kappa-1}$ otherwise a profitable and unobservable deviation would exist for a subset of sellers. ■

Proof Proposition 1.4.2.
Consider sellers in S^κ, $\kappa \geq 0$. Definition 1.4.1 requires separating strategies for every seller $i \in S^\kappa$—after all possible histories $h_i^{\kappa-1} \in H^{\kappa-1}$—on and off the equilibrium path.

Step 1. Let $S_H^\kappa, S_L^\kappa \neq \emptyset$. In equilibrium, separating strategies are:

$$\begin{aligned}
&\sigma_H^{*\kappa}(\theta_H|h_i^{\kappa-1}) = 1 && \sigma_L^{*\kappa}(\theta_L|h_i^{\kappa-1}) = 1 \\
&\sigma_B^*(A|\mathscr{I}_B(\theta_H, \cdot, H)) = 1 && \sigma_B^*(A|\mathscr{I}_B(\theta_H, \cdot, L)) = 0 \\
&\sigma_B^*(A|\mathscr{I}_B(\theta_L, \cdot, L)) = 1 && \sigma_B^*(A|\mathscr{I}_B(p, \cdot, \xi)) = 0 \text{ for } p \in (\theta_L, \theta_H)
\end{aligned}$$

for every $h_i^{\kappa-1} \in H^{\kappa-1}$.

Let $p_H^\kappa > p_L^\kappa$ be separating prices for H- and L-sellers in S^κ. As $S^\kappa \neq \emptyset$, on the equilibrium path sellers may stay on the market at least κ previous periods. In equilibrium, separation implies $\mathbb{E}_{\pi_B}[\theta|\mathscr{I}(p_H^\kappa, \cdot, \xi)] = \theta_H$ and $\mathbb{E}_{\pi_B}[\theta|\mathscr{I}(p_L^\kappa, \cdot, \xi)] = \theta_L$ for every $\xi \in \{H, L\}$.

If $\sigma_B(A|\mathscr{I}_B(p^{\kappa}_H,\cdot,\xi)) = \sigma_B(A|\mathscr{I}_B(p^{\kappa}_L,\cdot,\xi))$ for every $\xi \in \{H,L\}$, buyers accept p^{κ}_H and p^{κ}_L with equal probability. In turn L-sellers would deviate as $p^{\kappa}_H > p^{\kappa}_L$. Therefore, separation can occur only if buyers accept $(p^{\kappa}_H, p^{\kappa}_L)$ with different probabilities. By point 2. in Definition 3.3.1 the only plausible equilibrium strategy is: buyers accept p^{κ}_L for every $\xi \in \{H,L\}$ and p^{κ}_H only if $\xi = H$. This is possible only if $p^{\kappa}_H = \theta_H$ otherwise buyers would always accept. Sequential rationality requires L-sellers to ask the highest price accepted by buyers, i.e. $p^{\kappa}_L = \theta_L$.

Suppose *per contra* that separating behavioural strategies are not in pure strategies. From the argument in Step 1. of the proof of Lemma 1.4.1 sellers in S^{κ}_{λ} can mix: (i) between two prices p^{κ}_{λ} and $p^{\kappa}_{\lambda,2}$ both accepted with positive probability; (ii) between p^{κ}_{λ} and a price rejected with probability one.

If H-sellers play strategy (i), L-sellers would not post $p^{\kappa}_L = \theta_L$ as $p^{\kappa}_{H,2} \geq v_H$ is accepted with the same probability. Similarly, if L-sellers play strategy (i) and mix between θ_L and $p^{\kappa}_{L,2}$ they prefer not to post $p^{\kappa}_{L,2}$ but $\theta_H > p^{\kappa}_{L,2}$ as both prices are accepted only if $\xi = H$, contradicting the hypothesis that H-sellers play a separating strategy.

If sellers in S^{κ}_{λ} play strategy (ii) then at least one of these two indifference conditions hold:

$$\begin{aligned} \gamma(\theta_H - v_H) + (1-\gamma)V^{\kappa+1}_H(\sigma^*,\pi^*|h^{\kappa-1}_i) - c &= V^{\kappa+1}_H(\sigma^*,\pi^*|h^{\kappa-1}_i) - c \\ \theta_L - v_L - c &= V^{\kappa+1}_L(\sigma^*,\pi^*|h^{\kappa-1}_i) - c \end{aligned}$$

The first equation implies $V^{\kappa+1}_H(\sigma^*,\pi^*|h^{\kappa-1}_i) = \theta_H - v_H$ but in all possible equilibria it must be $V^{\kappa+1}_H(\sigma,\pi) \leq \theta_H - v_H - c$ otherwise buyers have to pay a price higher than θ_H, violating their individual rationality. The second equation cannot hold as well because a L-seller would get a higher payoff deviating to θ_H since:

$$(1-\gamma)(\theta_H - v_L) + \gamma V^{\kappa+1}_L(\sigma^*,\pi^*|h^{\kappa-1}_i) - c = (1-\gamma)(\theta_H - v_L) + \gamma(\theta_L - v_L) - c > \theta_L - v_L - c$$

Therefore, in equilibrium all L-sellers trade immediately while a share $1-\gamma$ of H-sellers in S^{κ}_H moves to period $\kappa+1$, i.e. $S^{\kappa+1}_H \neq \emptyset$ if H-sellers continue to participate in the market. Definition 1.4.1 requires separating behavioural strategies for every cohort S^{κ}, $\kappa \in \mathbb{N}_0$. Hence, the same separating behavioural strategy is played for every $\kappa \geq 0$. To support separation, out of equilibrium beliefs should be sufficiently negatively, say $\pi_B(p,\cdot,\xi) = 0$, for every price above θ_L but below θ_H.

Step 2. A separating behavioural strategy profile exists if and only if:

$$c \in \left[\frac{1-\gamma}{\gamma}(\theta_H - \theta_L), \gamma(\theta_H - v_H)\right]$$

By Step 1 and Lemma 1.8.2 it follows that $V_\lambda^\kappa(\sigma^*,\pi^*|h_i^{\kappa-1}) = V_\lambda(\sigma^*,\pi^*)$. Using Lemma 1.8.1 L-sellers in S_L^κ post θ_L and trade if and only if for every $h_i^{\kappa-1} \in H^{\kappa-1}$:

$$V_L(\sigma^*,\pi^*) = \theta_L - v_L - c \geq (1-\gamma)(\theta_H - v_L) + \gamma(\theta_L - v_L - c) - c$$

i.e. $c \geq \frac{1-\gamma}{\gamma}(\theta_H - \theta_L)$.

H-sellers' participate in the market until they trade if and only if:

$$V_H(\sigma^*,\pi^*) = \gamma(\theta_H - v_H) + (1-\gamma)V_H(\sigma^*,\pi^*) - c \geq 0$$

i.e. $V_H(\sigma^*,\pi^*) = \theta_H - v_H - \frac{c}{\gamma} \geq 0$ or $c \leq \gamma(\theta_H - v_H)$. ■

1.8.3 Pooling equilibria

Proof Lemma 1.4.1.
See Appendix B. ■

Lemma 1.8.3 *For c sufficiently small, under TMO there is no undefeated equilibrium (σ^*,π^*) in which $S_H^\kappa \neq \emptyset$, $S_L^\kappa \neq \emptyset$ and H- and L-sellers in S^κ only post prices rejected with probability one.*

Proof. By Lemma 1.4.1 if all sellers participate in the market the only admissible equilibrium strategies are: (i) H- and L-sellers post the same price and buyers accept only if $\xi = H$; or (ii) under TMO, H- and L-sellers only post prices rejected with probability one.

Let $K^N(\kappa) := \{j \geq \kappa : \text{sellers in } S^j \text{ play strategy (ii)}\}$. Obviously sellers in S_λ^κ, $\lambda \in \{H,L\}$, would not play strategy profile (ii) if they drop in the subsequent period $\kappa+1$. Therefore, there exists at least one future period in which they trade with positive probability, i.e. there exists at least one $l \geq \kappa$ such that sellers in S^l play behavioural strategy (i). Without loss of generality, consider j such that $j \in K^N(\kappa)$ and $j+1 \notin K^N(\kappa)$. Since no seller in S^j trades then $q^j = q^{j+1}$. Consider an alternative equilibrium $(\tilde{\sigma},\tilde{\pi})$ in which each seller $i \in S^j$ plays: $\tilde{\sigma}^l(p|h_i^{l-1}) = \sigma^{*^l}(p|h_i^{l-1})$ for every $l < j$ and $\forall h_i^{l-1} \in H^{l-1}$; and $\tilde{\sigma}^l(p|h_i^{l-1}) = \sigma^{*^{l+1}}(p|h_i^l)$ for every $l \geq j$ and $\forall h_i^{l-1} \in H^{l-1}$.[48]. Let $\tilde{\pi}(p,l,\xi) = \pi^*(p,l,\xi)$ for every $l < j$, and $\tilde{\pi}(p,l,\xi) = \pi^*(p,l+1,\xi)$ for every $l \geq j$. It is easy to observe that, if (σ^*,π^*) is an equilibrium assessment, $(\tilde{\sigma},\tilde{\pi})$ is also an equilibrium as it satisfies analogous no deviation conditions. However, $V_\lambda^l(\tilde{\sigma},\tilde{\pi}) > V_\lambda^l(\sigma^*,\pi^*)$ for every $l \leq j$ as they do not incur an extra cost c in period j. As a result, (σ^*,π^*) would be defeated by $(\tilde{\sigma},\tilde{\pi})$. ■

[48] The s-th element in h_i^{l-1} is equal to the s-th element in h^l for $s < j$ and equal to the $s+1$-th for $s \geq j$.

Proof Proposition 1.4.3.

1. By Lemma 1.4.1 and 1.8.3 the only admissible equilibrium strategy for $c \leq c^*$ is:

$$\sigma_H^{*^\kappa}(p^\kappa) = \sigma_L^{*^\kappa}(p^\kappa) = 1 \qquad \sigma_B^*(A|\mathscr{I}_B(p^\kappa,\kappa,H)) = 1 \qquad \sigma_B(A|\mathscr{I}_B(p^\kappa,\kappa,L)) = 0$$

Playing this strategy profile implies:

$$\left.\begin{array}{l} S_H^{\kappa+1} = (1-\gamma)S_H^\kappa \\ S_L^{\kappa+1} = \gamma S_L^\kappa \end{array}\right\} \tfrac{S^{\kappa+1}}{S^\kappa} = (1-\gamma)\tfrac{S_H^\kappa}{S^\kappa} + \gamma\tfrac{S_L^\kappa}{S^\kappa} = (1-\gamma) + (2\gamma-1)(1-q^\kappa)$$

Therefore:

$$q^{\kappa+1} = \frac{S_H^{\kappa+1}}{S^{\kappa+1}} = (1-\gamma)\frac{S_H^\kappa}{S^\kappa}\frac{S^\kappa}{S^{\kappa+1}} = \frac{(1-\gamma)q^\kappa}{(1-\gamma)+(2\gamma-1)(1-q^\kappa)} := g(q^\kappa) := g^\kappa(q^0)$$

As $q^{\kappa+1}$ does not depend on p^κ and buyers cannot observe previously posted prices, future continuation values $V_\lambda^j(\sigma^*,\pi^*)$, $j > \kappa$, do not depend on p^κ. As a result, $p^\kappa = \mathbb{E}_{\pi_B}[\theta|\mathscr{I}(p^\kappa,\kappa,H)]$ supports the unique undefeated equilibrium (σ^*,π^*) as both types of seller get the highest possible payoff in the class of admissible equilibria (see Lemma 1.4.1). The undefeated equilibrium is unique because prices $\{p^\kappa\}_{\kappa\in\mathbb{N}_0}$ are unique.

Let $V_\lambda^\kappa(q) := V_\lambda^\kappa(\sigma^*,\pi^*)$ be the continuation value in (σ^*,π^*) for a seller $i \in S_\lambda^\kappa$ when $q^\kappa = q$. Let $\kappa^*(q^0)+1$ be the maximum number of periods on the market for a H-seller. If $\kappa^*(q^0) = 0$ he participates just for one period, while if $\kappa^*(q^0) = +\infty$ he participates until he trades. For every seller $i \in S^\kappa$, $\kappa \leq \kappa^*(q^0)$, the maximization problem can be rewritten as follows:

$$V_\lambda^\kappa(q^\kappa) = \mathbb{P}_\lambda(H)(p^\kappa - v_\lambda) + [1-\mathbb{P}_\lambda(H)]V_\lambda^{\kappa+1}(q^{\kappa+1}) - c$$

As q^κ is decreasing in κ then p^κ is also decreasing in κ as $\mathbb{E}_{\pi_B}[\theta|\mathscr{I}_B(p^\kappa,\kappa,H)]$ is monotonically increasing in q^κ. Hence, $V_\lambda^\kappa(q^\kappa)$ is decreasing in q^κ and $V_\lambda^{\kappa+1}(q^{\kappa+1}) < V_\lambda^\kappa(q^\kappa)$ as $q^{\kappa+1} < q^\kappa$. As a result, H-sellers do not find profitable to postpone trade to a future period.

Market participation requires $V_\lambda^\kappa(q^\kappa) \geq 0$. Let's consider the following bounds on $V_H^\kappa(q^\kappa)$:

Upper bound: $U_H^\kappa(q^\kappa) = \gamma(p^\kappa - v_H) + (1-\gamma)U_H^\kappa(q^\kappa) - c \Rightarrow U_H^\kappa(q^\kappa) = p^\kappa - v_H - \frac{c}{\gamma}$
Lower bound: $L_H^\kappa(q^\kappa) = \gamma(p^\kappa - v_H) - c$

$U_H^\kappa(q^\kappa)$, $V_H^\kappa(q^\kappa)$ and $L_H^\kappa(q^\kappa)$ are monotonically increasing in q^κ and $U_H^\kappa(q^\kappa) \geq V_H^\kappa(q^\kappa) \geq L_H^\kappa(q^\kappa)$. Notice that $U_H^\kappa(q^\kappa) = \frac{L_H^\kappa(q^\kappa)}{\gamma}$ so H-sellers exit the market whenever $L_H^\kappa(q^\kappa) \leq 0$.

Let $q_c^O := \underset{q\in(0,1)}{\arg}\ L_H^\kappa(q) = 0$. Then:

$$L_H^\kappa(q_c^O) = \gamma(\mathbb{E}_{\pi_B}[\theta|\mathscr{I}_B(p^\kappa,\kappa,H)] - v_H) - c = \gamma\left(\frac{q_c^O\gamma\theta_H + (1-q_c^O)(1-\gamma)\theta_L}{q_c^O\gamma + (1-\gamma)(1-q_c^O)} - v_H\right) - c = 0$$

Solving for q_c^O

$$q_c^O = \frac{(1-\gamma)(v_H + \frac{c}{\gamma} - \theta_L)}{\gamma(\theta_H - v_H - \frac{c}{\gamma}) + (1-\gamma)(v_H + \frac{c}{\gamma} - \theta_L)}$$

H-sellers participate only for a finite number of periods:

$$\kappa^*(q^0) = \max\left\{\kappa \in \mathbb{N}_0 : g^\kappa(q^0) \geq q_c^O\right\}$$

since $q^{\kappa+1} = g^\kappa(q^0)$ is strictly decreasing in κ.

2. H-sellers do not participate in economies $\mathscr{E}(\gamma, q^0)$ such that $q^0 < q_c^O$. L-sellers always trade because there are gains from trade ($\theta_L > v_L$), and the only possible sequential best response is to post θ_L which buyers accept with probability one.[49] ■

Proof Proposition 1.4.4.

1. By Lemma 1.4.1, for $c \leq c^*$ sellers' equilibrium strategy are $\sigma_\lambda(p) = \sigma_\lambda^\kappa(p) = 1$ for every $\lambda \in \{H, L\}$ and $\kappa \in \mathbb{N}_0$. Price p is accepted only if $\xi = H$.

In equilibrium, strategies do not depend on time and, by Lemma 1.4.1, the price p is posted by all cohorts of sellers. Thus, this outcome is possible only if the economy is in a stationary state, i.e. if the measure of sellers S, say $\bar{S}$, and the fraction of H-sellers, say $\bar{q} = \frac{\bar{S}_H}{\bar{S}}$, are constant over time. In turn, in every period an equal measure of each type of seller must enter and exit the market:

$$\begin{cases} q^0 = \bar{S}\gamma\bar{q} \\ (1-q^0) = \bar{S}(1-\gamma)(1-\bar{q}) \end{cases} \Rightarrow \begin{cases} \bar{S} = \frac{\gamma - q^0(2\gamma-1)}{\gamma(1-\gamma)} \\ \bar{q} = \frac{q^0(1-\gamma)}{\gamma - q^0(2\gamma-1)} \end{cases}$$

Notice that $\bar{S}$ and $\bar{q}$ do not depend on c or p. A similar argument to the one presented in the proof of Proposition 1.4.3 ensures that the undefeated equilibrium is unique and sellers post a price $\bar{p}$ equal to buyers' posterior valuation when $\xi = H$.

$$\bar{p} = \mathbb{E}_{\pi_B}[\theta|\mathscr{I}_B(\bar{p}, H)] = \frac{\gamma\bar{q}\theta_H + (1-\gamma)(1-\bar{q})\theta_L}{\gamma\bar{q} + (1-\gamma)(1-\bar{q})} = \frac{\theta_H + \frac{(1-\gamma)}{\gamma}\frac{(1-\bar{q})}{\bar{q}}\theta_L}{1 + \frac{(1-\gamma)}{\gamma}\frac{(1-\bar{q})}{\bar{q}}}$$

[49] The price θ_L must be always accepted with probability one; otherwise, L-sellers would deviate and post $\theta_L - \varepsilon$, for ε arbitrary small, which buyers would always accept. Indeed, under any belief, their minimum valuation for the good is θ_L.

Observe that $\frac{(1-\gamma)}{\gamma}\frac{(1-\bar{q})}{\bar{q}} = \frac{1-q^0}{q^0}$, so $\bar{p} = \mathbb{E}_{\pi_B}[\theta|\mathscr{I}_B(\bar{p},H)] = q^0\theta_H + (1-q^0)\theta_L$.

Every cohort of sellers S^κ posts price $\bar{p}$ and previous search costs are sunk, hence continuation values are constant $\forall\kappa$, i.e. $V^\kappa_\lambda(\sigma^*,\pi^*) = V_\lambda(\sigma^*,\pi^*)$. H-sellers participate in the market until they trade because their forward looking decision problem is unchanged. The individual rationality constraint for H-sellers is:

$$V_H(\sigma^*,\pi^*) = \gamma(\bar{p}-v_H) + (1-\gamma)V_H(\sigma^*,\pi^*) - c \geq 0$$

i.e.

$$V_H(\sigma^*,\pi^*) = \bar{p} - v_H - \frac{c}{\gamma} = q^0\theta_H + (1-q^0)\theta_L - v_H - \frac{c}{\gamma} \geq 0.$$

This is satisfied only if $q^0 \geq q^N_c = \frac{v_H - \theta_L + \frac{c}{\gamma}}{\theta_H - \theta_L}$.

2. See point 2. in the proof of Proposition 1.4.3. ■

1.8.4 Welfare analysis

Proof Proposition 1.5.1.

It follows directly from Definition 1.5.1 and Propositions 1.4.3 - 1.4.4. ■

Proof Proposition 1.5.2.

I construct an equilibrium assessment (σ,π) such that:

1. For every $\kappa \in \mathbb{N}_0$ H-sellers always post the price $p^\kappa_H = p_H = \theta_L + \frac{\gamma}{1-\gamma}c$ while L-sellers post p_H with probability $\frac{1-\gamma}{\gamma}$ and θ_L with probability $\frac{2\gamma-1}{\gamma}$.

2. Buyers always accept θ_L while they accept p_H only if $\xi = H$.

Step 1. Equilibrium characterization

In a semi-separating equilibrium sellers evolve across cohorts S^κ according to:

$$\left.\begin{aligned} S^{\kappa+1}_H &= (1-\gamma)S^\kappa_H \\ S^{\kappa+1}_L &= \sigma^\kappa_L(p_H)\gamma S^\kappa_L \end{aligned}\right\} \frac{S^{\kappa+1}}{S^\kappa} = (1-\gamma)\frac{S^\kappa_H}{S^\kappa} + \sigma^\kappa_L(p_H)\gamma\frac{S^\kappa_L}{S^\kappa} = (1-\gamma) + [\gamma(1+\sigma^\kappa_L(p_H)] - 1](1-q^\kappa)$$

Hence:

$$q^{\kappa+1} = \frac{S^{\kappa+1}_H}{S^{\kappa+1}} = (1-\gamma)\frac{S^\kappa_H}{S^\kappa}\frac{S^\kappa}{S^{\kappa+1}} = \frac{(1-\gamma)}{(1-\gamma) + [\gamma(1+\sigma^\kappa_L(p_H)) - 1](1-q^\kappa)}q^\kappa$$

If $q^{\kappa+1} = q^\kappa$ $\forall\kappa \in \mathbb{N}_0$ L-sellers' strategy satisfies $[\gamma(1+\sigma^\kappa_L(p_H)] - 1] = 0$, i.e. $\sigma^k_L(p_H) = \frac{1-\gamma}{\gamma}$. To play a mixed behavioural strategy L-sellers are indifferent between posting p_H or θ_L. Since $q^\kappa = q^0$ for every $\kappa \in \mathbb{N}_0$, time on market observability is irrelevant because buyers' prior probability to match with a H-seller is constant across cohorts. Buyers' optimal strategy does not change across cohorts S^κ and—in turn—sellers' continuation values do not depend on κ, i.e. $V_\lambda^\kappa(q^\kappa) = V_\lambda(q^0)$[50] $\forall \kappa \in \mathbb{N}_0$. The indifference condition for L-sellers is:

$$V_L(q^0) = (1-\gamma)(p_H - v_L) + \gamma V_L(q^0) - c = \theta_L - v_L - c$$

or

$$V_L(q^0) = (1-\gamma)(p_H - v_L) + \gamma[\theta_L - v_L - c] - c = \theta_L - v_L - c$$

The last expression implies $p_H = \theta_L + \frac{\gamma}{1-\gamma}c$. Substituting this price into:

$$V_H(q^0) = \gamma(p_H - v_H) + (1-\gamma)V_H(q^0) - c$$

I get:

$$V_H(q^0) = p_H - v_H - \frac{c}{\gamma} = \theta_L - v_H + \frac{\gamma^2+\gamma-1}{\gamma(1-\gamma)}c$$

H-sellers individual rationality constraint is satisfied only if $c \geq \frac{\gamma(1-\gamma)}{\gamma^2+\gamma-1}(v_H - \theta_L)$.

Buyers follow their equilibrium strategy if and only if:

$$\mathbb{E}_{\pi_B}[\theta|\mathscr{I}_B(p_H,\kappa,H)] \geq p_H \geq \mathbb{E}_{\pi_B}[\theta|\mathscr{I}_B(p_H,\kappa,L)]$$

i.e.

$$\frac{q^0\gamma\theta_H + (1-q^0)\frac{(1-\gamma)^2}{\gamma}\theta_L}{q^0\gamma + (1-q^0)\frac{(1-\gamma)^2}{\gamma}} \geq \theta_L + \frac{\gamma}{1-\gamma}c \geq q^0\theta_H + (1-q^0)\theta_L$$

This set of inequalities can be rewritten as:

$$q_{SSU} := \frac{\frac{\gamma}{1-\gamma}c}{\theta_H - \theta_L} \geq q^0 \geq \frac{1-\gamma}{\gamma\left[\frac{\theta_H-\theta_L}{c}\right] - \frac{2\gamma-1}{1-\gamma}} := q_{SS}$$

Notice that $\gamma\left[\frac{\theta_H-\theta_L}{c}\right] - \frac{2\gamma-1}{1-\gamma} > 0$ since $\frac{\theta_H-\theta_L}{c} \geq \frac{\gamma}{1-\gamma} > \frac{2\gamma-1}{\gamma(1-\gamma)}$.

Step 2. There exist a set of economies $\mathscr{E}(\gamma, q^0)$ that support the equilibrium in Step 1. but not the pooling equilibria in Propositions 1.4.3 and 1.4.4.

It is sufficient to find a set of parameters $(\theta_H, \theta_L, v_H, v_L, \gamma, c)$ such that $q^{SS} < q_0^O$, i.e.:

$$\frac{1-\gamma}{\gamma\left[\frac{\theta_H-\theta_L}{c}\right] - \frac{2\gamma-1}{1-\gamma}} < \frac{(1-\gamma)(v_H-\theta_L)}{\gamma(\theta_H - v_H) + (1-\gamma)(v_H - \theta_L)}$$

and c is such that the semi-separating equilibrium exists, i.e. $c \geq \frac{\gamma(1-\gamma)}{\gamma^2+\gamma-1}(v_H - \theta_L)$.

[50] I adopt the notation already used in the proof of Proposition 1.4.3.

As $v_H - \theta_L > 0$, the signal precision γ must be sufficiently high to have $\frac{\gamma^2+\gamma-1}{\gamma(1-\gamma)} > 0$, i.e. $\gamma > \frac{\sqrt{5}-1}{2}$. For a given set of parameters $(\theta_H, \theta_L, v_H, v_L, \gamma)$, the lower bound q^{SS} is increasing in c so its value is minimum for $c = \frac{\gamma(1-\gamma)}{\gamma^2+\gamma-1}(v_H - \theta_L)$. Then, $q^{SS} < q_0^O$ requires:

$$\frac{1-\gamma}{\gamma\left[\frac{\theta_H-\theta_L}{c}\right] - \frac{2\gamma-1}{1-\gamma}} = \frac{(1-\gamma)^2(v_H-\theta_L)}{(\gamma^2+\gamma-1)\theta_H + \gamma(1-\gamma)\theta_L - (2\gamma-1)v_H} < \frac{(1-\gamma)(v_H-\theta_L)}{\gamma(\theta_H - v_H) + (1-\gamma)(v_H - \theta_L)}$$

After some tedious calculations, the above inequality is equivalent to:

$$2\frac{\theta_H - v_H}{v_H - \theta_L}\gamma^2 + \gamma - \frac{\theta_H - \theta_L}{v_H - \theta_L} > 0$$

Using the standard quadratic formula:

$$\gamma > \hat{\gamma} = \frac{-1 + \sqrt{1 - 8\frac{(\theta_H - v_H)(\theta_H - \theta_L)}{(v_H - \theta_L)^2}}}{4\frac{\theta_H - v_H}{v_H - \theta_L}}$$

The discriminant $\Delta \equiv 1 - 8\frac{(\theta_H - v_H)(\theta_H - \theta_L)}{(v_H - \theta_L)^2}$ is strictly positive if $\frac{(\theta_H - v_H)(\theta_H - \theta_L)}{(v_H - \theta_L)^2} < \frac{1}{8}$ (for example: $\theta_H = 1$, $v_H = 0.98$ and $\theta_L = 0.8$). If the latter inequality holds, then $\Delta < 1$ and $\hat{\gamma} < 1$ because:

$$\frac{-1 + \sqrt{\Delta}}{4\frac{\theta_H - v_H}{v_H - \theta_L}} < 1 \Leftrightarrow \sqrt{\Delta} - 4\frac{\theta_H - v_H}{v_H - \theta_L} < 1$$

Therefore, if $\frac{(\theta_H - v_H)(\theta_H - \theta_L)}{(v_H - \theta_L)^2} < \frac{1}{8}$ for every $\gamma > \max(\hat{\gamma}, \frac{\sqrt{5}-1}{2})$ there exists an economy such that $q^{SS} < q_0^O$. ∎

1.8.5 Market design

Proof Proposition 1.6.1.

Under TMO the set of feasible transfers $(\tau^\kappa, r^\kappa)_{\kappa=0}^{+\infty}$ can vary across sellers belonging to different cohorts S^κ. I denote with τ^κ a transfer from a participating seller in S^κ to the designer and with r^κ a transfer from the designer to a seller, if he trades after κ previous periods in the market.[51] For a constellation $(\theta_H, v_H, \theta_L, v_L, \gamma)$, let $(\tau^{*\kappa}, r^{*\kappa})_{\kappa=0}^{+\infty}$ be the efficient mechanism existing for the lowest possible q^0. Denote this minimum value with q^*.

To implement an equilibrium with informationally efficient prices, the mechanism

[51] For example $\tau^\kappa > 0$ is the amount that a seller in S^κ has to pay to the designer in order to participate in the market. On the contrary $r^\kappa > 0$ is the positive transfer that a seller receives if he trades after κ previous periods in the market. In this vein, I often call τ^κ a market participation tax and r^κ a tax rebate that a seller receives upon trade.

has to implement a separating equilibrium. By Step 1 and 2 of the proof of Proposition 1.4.2, in equilibrium, L-sellers trade immediately while H-sellers trade once they match with a buyer who receives a high signal. By the appropriate law of large numbers, a system of transfers is budget balanced (henceforth BB) on the equilibrium path if and only if:

$$(1-q^0)(r^0-\tau^0) \leq q^0 \sum_{\kappa=0}^{+\infty} (1-\gamma)^\kappa (\tau^\kappa - \gamma r^\kappa)$$

Since utility is quasi-linear in transfers and agents are risk neutral, if the BB constraint holds with equality a feasible mechanism does not change the total trade surplus between buyers and sellers, i.e. $q^0(\theta_H - v_H) + (1-q^0)(\theta_L - v_L)$. Moreover, for $c = 0$ delaying trade does not reduce total surplus. In a separating equilibrium assessment, say (σ^*, π^*), sellers of type λ trade at price θ_λ and buyers' expected payoff is zero. Therefore, if BB holds with equality a mechanism implementing a separating equilibrium (σ^*, π^*) satisfies:

$$q^0 V_H^0(\sigma^*, \pi^*) + (1-q^0) V_L^0(\sigma^* \pi^*) = q^0(\theta_H - v_H) + (1-q^0)(\theta_L - v_L)$$

where the LHS is sellers' equilibrium payoffs while the RHS is total trade surplus.

If BB is slack it is possible to have an efficient mechanism for every q^0, i.e. $q^* = 0$. If $q^* > 0$, BB holds with equality, otherwise the market designer may use his profit to relax sellers' individual rationality constraints. In the remainder of the proof I consider a binding BB.

Step 1. If $V_H^0(\sigma^*, \pi^*) = 0$ then $V_H^\kappa(\sigma^*, \pi^*) = 0$ for every $\kappa \in \mathbb{N}_0$.

If $V_H^0(\sigma^*, \pi^*) = 0$ then $(1-q^0)V_L^0(\sigma^*, \pi^*) = q^0(\theta_H - v_H) + (1-q^0)(\theta_L - v_L)$, i.e. L-sellers get all trade surplus available in the economy. If $V_H^\kappa(\sigma^*, \pi^*) < 0$ for some κ then H-sellers in S_H^κ would exit the market, and the resulting allocation would not be efficient. If $V_H^\kappa(\sigma^*, \pi^*) > 0$ for some $\kappa \geq 1$, then H-sellers in cohort S^κ enjoy a strictly positive payoff only if their matched buyers get a negative expected payoff, contradicting the optimality of buyers' strategy.

Step 2. If $q^* > 0$ then $V_H^0(\sigma^*, \pi^*) = 0$.

As $(\tau^{*\kappa}, r^{*\kappa})_{\kappa=0}^{+\infty}$ supports a separating equilibrium H-sellers' equilibrium payoff is:[52]

$$V_H^0(q^0) = \theta_H - v_H - \sum_{\kappa=0}^{\infty} (1-\gamma)^\kappa (\tau^{*\kappa} - \gamma r^{*\kappa}) = \theta_H - v_H - \frac{1-q^0}{q^0}(r^{*0} - \tau^{*0})$$

where the last equality follows from the BB constraint. Observe that τ^{*0} only affects

[52] Similarly to the proof of Proposition 1.4.3 I stress the importance of q^0 and I expand the notation using $V_\lambda^\kappa(q^0)$ rather than $V_\lambda^\kappa(\sigma^*, \pi^*)$.

the initial market participation decision for sellers in cohorts $\kappa = 0$, but it does not change the future incentive compatibility constraints. Rearranging the BB constraint:

$$\tau^{*^0} = (1-q^0)r^{*^0} - q^0[\gamma r^{*^0} + \sum_{\kappa=1}^{\infty}(1-\gamma)^\kappa(\tau^{*^\kappa} - \gamma r^{*^\kappa})]$$

For different values of q^0 we can leave unchanged r^{*^0} and $(\tau^{*^\kappa}, r^{*^\kappa})_{\kappa=1}^{\infty}$ and achieve BB through τ^0. Substituting τ^{*^0} in sellers' expected payoff:[53]

$$\begin{aligned}
V_H^0(q^0) = &\quad \theta_H - v_H - \tfrac{1-q^0}{q^0}(r^{*^0} - \tau^{*^0}) \\
= &\quad \theta_H - v_H - (1-q^0)[(1+\gamma)r^{*^0} + \sum_{\kappa=1}^{\infty}(1-\gamma)^\kappa(\tau^{*^\kappa} - \gamma r^{*^\kappa})] \\
V_L^0(q^0) = &\quad \theta_L - v_L + \tfrac{q^0}{1-q^0}[\theta_H - v_H - V_H^0(\sigma^*)] \\
= &\quad \theta_L - v_L + q^0[(1+\gamma)r^{*^0} + \sum_{\kappa=1}^{\infty}(1-\gamma)^\kappa(\tau^{*^\kappa} - \gamma r^{*^\kappa})]
\end{aligned}$$

Suppose on the contrary that $V_H^0(q^*) > 0$ with $q^* > 0$. If $[(1+\gamma)r^{*^0} + \sum_{\kappa=1}^{\infty}(1-\gamma)^\kappa(\tau^{*^\kappa} - \gamma r^{*^\kappa})] > 0$ then it is easy to observe that there exists a sufficiently small $\varepsilon > 0$ such that for every $q^0 \geq q^* - \varepsilon$ a market intervention that only adjusts τ^0 to preserve budget balance achieves $V_H^\kappa(q^0) \geq 0$ and $V_L^\kappa(q^0) \geq 0$ for every $\kappa \in \mathbb{N}_0$, contradicting the definition of q^*. If instead $[(1+\gamma)r^{*^0} + \sum_{\kappa=1}^{\infty}(1-\gamma)^\kappa(\tau^{*^\kappa} - \gamma r^{*^\kappa})] < 0$ then for every $q^0 < q^*$ we can adjust τ^0 such that $V_H^0(q^0) > V_H^0(q^*)$ and $V_L^0(q^0) > V_L^0(q^*) \geq 0$, contradicting again the definition of q^*.

Step 3. For every $(\theta_H, v_H, \theta_L, v_L)$ there exists γ^* such that for $\gamma \in (\frac{1}{2}, \gamma^*)$ it is $q^* > 0$.

Let $(\tau_\gamma^{*^\kappa}, r_\gamma^{*^\kappa})_{\kappa=0}^{\infty}$ be a mechanism that implements a separating equilibrium in an economy with signal precision γ. Sellers in S_L^0 post θ_L only if:

$$\theta_L - v_L + r_\gamma^{0^*} - \tau_\gamma^{0^*} \geq \theta_H - v_L - \sum_{\kappa=0}^{\infty}\gamma^\kappa[\tau_\gamma^{\kappa^*} - (1-\gamma)r_\gamma^{\kappa^*}] \tag{1.7}$$

From the BB constraint I can substitute $r_\gamma^{0^*} - \tau_\gamma^{0^*}$ with $\frac{q^0}{1-q^0}\sum_{\kappa=0}^{\infty}(1-\gamma)^\kappa(\tau_\gamma^{*^\kappa} - \gamma r_\gamma^{*^\kappa})]$ to get the no-deviation condition:

$$\sum_{\kappa=0}^{\infty}\gamma^\kappa[\tau_\gamma^{\kappa^*} - (1-\gamma)r_\gamma^{\kappa^*}] + \frac{q^0}{1-q^0}\sum_{\kappa=0}^{\infty}(1-\gamma)^\kappa(\tau_\gamma^{\kappa^*} - \gamma r_\gamma^{\kappa^*})] \geq \theta_H - \theta_L$$

[53] Notice that I use the trade surplus equation to rewrite L-sellers expected payoff.

Notice that for every $\varepsilon > 0$ it is possible to find γ_ε^* such that for every $\gamma \in (\frac{1}{2}, \gamma_\varepsilon^*)$:

$$|\Delta_\gamma| := \left| \sum_{\kappa=0}^{\infty} \gamma^\kappa [\tau_\gamma^{\kappa^*} - (1-\gamma) r_\gamma^{\kappa^*}] - \sum_{\kappa=0}^{\infty} (1-\gamma)^\kappa (\tau_\gamma^{*^\kappa} - \gamma r_\gamma^{*^\kappa})] \right| < \varepsilon$$

where $\Delta_\gamma \to 0$ for $\gamma \to \frac{1}{2}$.
Adding $\sum_{\kappa=0}^{\infty} (1-\gamma)^\kappa (\tau_\gamma^{*^\kappa} - \gamma r_\gamma^{*^\kappa})]$ to both sides of equation (1.7) and simplifying:

$$\sum_{\kappa=0}^{\infty} (1-\gamma)^\kappa (\tau_\gamma^{*^\kappa} - \gamma r_\gamma^{*^\kappa})] \geq (1-q^0)(\theta_H - \theta_L - \Delta_\gamma)$$

Now observe that the individual rationality constraint for H-sellers requires:

$$V_H^0(\sigma^*, \pi^*) = \theta_H - v_H - \sum_{\kappa=0}^{\infty} (1-\gamma)^\kappa (\tau_\gamma^{*^\kappa} - \gamma r_\gamma^{*^\kappa})] \geq 0$$

However for γ sufficiently close to $\frac{1}{2}$ (i.e. $\Delta_\gamma \to 0$) and q^0 small enough we have:

$$\theta_H - v_H - \sum_{\kappa=0}^{\infty} (1-\gamma)^\kappa (\tau_\gamma^{*^\kappa} - \gamma r_\gamma^{*^\kappa})] \leq \theta_H - v_H - (1-q^0)(\theta_H - \theta_L - \Delta_\gamma) < 0$$

as $v_H > \theta_L$. Therefore, no feasible system of transfers can implement an efficient allocation for values of q^0 and γ sufficiently close to zero and $\frac{1}{2}$, respectively.

Step 4. No separating equilibrium exists for

$$q^0 < \max\left\{0, 1 - \left(\frac{\gamma}{1-\gamma}\right)^2 \frac{\theta_H - v_H}{\theta_H - \theta_L}\right\} = q^*$$

A first best-allocation is implemented with a market participation tax $\tau^\kappa = \tau^* = \frac{1-\gamma}{\gamma}(\theta_H - \theta_L)$ *and, once a seller trades, a tax rebate* $r^\kappa = r^* = \tau^* \left(1 + \frac{1-\gamma}{\gamma} q^*\right)$.

By Step 1 and 2 if $q^* > 0$ we must have $V_H^\kappa(q^*) = 0$ for every $\kappa \in \mathbb{N}_0$, hence:

$$V_H^\kappa(q^*) = \gamma(\theta_H - v_H + r^{*^\kappa}) - \tau^{*^\kappa} = 0 \Rightarrow \tau^{*^\kappa} - \gamma r^{*^\kappa} = \gamma(\theta_H - v_H)$$

Substituting this expression into the BB constraint:

$$(1-q^0)(r^{*^0} - \tau^{*^0}) = q^0 \sum_{\kappa=0}^{+\infty} (1-\gamma)^\kappa (\tau^{*^\kappa} - \gamma r^{*^\kappa}) = q^0 \sum_{\kappa=0}^{+\infty} (1-\gamma)^\kappa \gamma(\theta_H - v_H) = q^0(\theta_H - v_H)$$

Hence $r^{*^0} - \tau^{*^0} = \frac{q^0}{1-q^0}(\theta_H - v_H)$. Combining with $\tau^{*^0} - \gamma r^{*^0} = \gamma(\theta_H - v_H)$ I get:

$$r^{*^0} = \frac{\theta_H - v_H}{1-\gamma}\frac{\gamma + q^0(1-\gamma)}{1-q^0} \qquad \tau^{*^0} = \frac{\gamma}{1-\gamma}\frac{\theta_H - v_H}{1-q^0} \tag{1.8}$$

Using the one-shot deviation property, L-sellers no deviation condition has to satisfy:

$$\theta_L - v_L + r^{*^\kappa} - \tau^{*^\kappa} \geq (1-\gamma)(\theta_H - v_L + r^{*^\kappa}) + \gamma(\theta_L - v_L + r^{*^{\kappa+1}} - \tau^{*^{\kappa+1}}) - \tau^{*^\kappa}$$

$$r^{*^\kappa} + \tau^{*^{\kappa+1}} - r^{*^{\kappa+1}} \geq \frac{1-\gamma}{\gamma}(\theta_H - \theta_L) \tag{1.9}$$

As $\tau^{*^{\kappa+1}} = \gamma(\theta_H - v_H + r^{*^{\kappa+1}})$ for every $\kappa \in \mathbb{N}_0$, substituting this value in equation (1.9):

$$r^{*^{\kappa+1}} \leq \frac{r^{*^\kappa}}{1-\gamma} + \frac{\gamma}{1-\gamma}(\theta_H - v_H) - \frac{\theta_H - \theta_L}{\gamma}$$

Observe that sellers in S_H^κ prefer not to postpone trade to period $\kappa+1$ only if:

$$\gamma(\theta_H - v_H + r^{*^\kappa}) - \tau^{*^\kappa} \geq V_H^{\kappa+1}(\sigma^*, q^*) - \tau^{*^\kappa} = -\tau^{*^\kappa} \Rightarrow \; r^{*^\kappa} \geq -(\theta_H - v_H)$$

where $V_H^{\kappa+1}(\sigma^*, q^*) = 0$ follows from Step 1. Therefore, for every $\kappa \geq 1$, r^{*^κ} satisfies:

$$-(\theta_H - v_H) \leq r^{*^\kappa} \leq \frac{r^{*^{\kappa-1}}}{1-\gamma} + \frac{\gamma}{1-\gamma}(\theta_H - v_H) - \frac{\theta_H - \theta_L}{\gamma} \tag{1.10}$$

In order to satisfy the LHS of equation (1.10) consider the RHS inequality binding. In this case, r^{*^κ} satisfies a first-order difference equation with solution:

$$r^{*^\kappa} = \left(\frac{1}{1-\gamma}\right)^\kappa \left[r^{*^0} + (\theta_H - v_H) - \left(\frac{1-\gamma}{\gamma^2}\right)(\theta_H - \theta_L)\right] - \left[(\theta_H - v_H) - \left(\frac{1-\gamma}{\gamma^2}\right)(\theta_H - \theta_L)\right]$$

Substituting r^{*^0} from equation (1.8) and rearranging:

$$r^{*^\kappa} = \left(\frac{1}{1-\gamma}\right)^\kappa \left[\frac{\theta_H - v_H}{(1-\gamma)(1-q^0)} - \left(\frac{1-\gamma}{\gamma^2}\right)(\theta_H - \theta_L)\right] - \left[(\theta_H - v_H) - \left(\frac{1-\gamma}{\gamma^2}\right)(\theta_H - \theta_L)\right]$$

As $\frac{1}{1-\gamma} > 1$, the solution is not explosive towards $-\infty$—violating equation (1.10)—only if the first term in squared brackets is non-negative. Simplifying the expression:

$$q^0 \geq 1 - \left(\frac{\gamma}{1-\gamma}\right)^2 \frac{\theta_H - v_H}{\theta_H - \theta_L}$$

Let $q^* := \max\left\{0, 1 - \left(\frac{\gamma}{1-\gamma}\right)^2 \frac{\theta_H - v_H}{\theta_H - \theta_L}\right\}$. All the previous inequalities constraints are binding when $q^0 = q^* > 0$. Then:

$$\tau^\kappa = \tau^* = \frac{1-\gamma}{\gamma}(\theta_H - \theta_L) \qquad r^\kappa = r^* = \tau^*\left(1 + \frac{1-\gamma}{\gamma} q^*\right)$$

■

1.9 Appendix B

Outline of the proof for Lemma 1.4.1.
I show that no other behavioural strategy is admissible except for the ones in Lemma 1.4.1. To prove this result, I first obtain a bound on the difference between continuation values, i.e. $V_L^\kappa - V_H^\kappa$. This preliminary result is obtained in three Lemmata. First, I show that, for c sufficiently small, there is a common price path played by H- and L-sellers (Lemma 1.9.1 and 1.9.2); then, I derive an expression for V_λ^κ and a bound on their difference (Lemma 1.9.3).

Lemma 1.9.1 *For c sufficiently small, no equilibrium path has L-sellers in $S_L^\kappa \neq \emptyset$ trade with positive probability, and H-sellers in $S_H^\kappa \neq \emptyset$ only post prices rejected with probability one.*

Proof Lemma 1.9.1.
Consider an equilibrium (σ^*, π^*) and sellers $i \in S_H^\kappa$ and $l \in S_L^\kappa$ with histories $h_i^{\kappa-1}$ and $h_l^{\kappa-1}$. Let $\bar{\kappa} := \arg\max_{m \in \mathbb{N}_0} \left\{ \exists \{p^j\}_{j=\kappa}^m : \sigma_H^{*j}(p^j | h_i^{j-1}) > 0, h_i^{j-1} = \{h_i^{\kappa-1}, p^\kappa, ..., p^{j-1}\} \right\}$.[54] Let $\tilde{h}_i^{\bar{\kappa}-1}$ be the history resulting from $h_i^{\kappa-1}$ and the sequence of prices $\{p^j\}_{j=\kappa}^{\bar{\kappa}-1}$ up to period $\bar{\kappa}$. Denote with $\tilde{h}_i^{j-1}$, $j \leq \bar{\kappa}$, a sub-history of $\tilde{h}_i^{\bar{\kappa}-1}$. If $\bar{\kappa} < +\infty$ the seller drops out after $\bar{\kappa}+1$ periods.

Suppose, per contra, that sellers in S^κ, $\kappa \leq \bar{\kappa} < \infty$, play the behavioural strategy described in the Lemma statement. First, sequential rationality implies $\kappa < \bar{\kappa}$: posting a price rejected with probability one in period $\bar{\kappa}$ before dropping out in period $\bar{\kappa}+1$ is not optimal, since dropping out after $\bar{\kappa}$ previous periods saves one search cost c. As a consequence, price $p^{\bar{\kappa}}$ is accepted with positive probability. Similarly, if $\bar{\kappa} = +\infty$, there is at least a $\kappa' > \kappa$ such that H-sellers in $S^{\kappa'}$ post a price $p^{\kappa'}$ accepted with positive probability.

By assumption, L-sellers trade with positive probability. Sequential rationality implies they post θ_L. In turn, seller $i \in S_L^\kappa$ finds profitable to trade at θ_L rather than postponing trade only if:

$$\theta_L - v_L - c \geq V_L^{\kappa+1}(\sigma^*, \pi^* | (h_l^{\kappa-1}, p^\kappa \times R)) - c \tag{1.11}$$

[54] If $\bar{\kappa} = +\infty$ there is at least one history in which a high quality seller participates in the market forever.

A possible deviation for a L-seller $l \in S_L^{\kappa+1}$ is to imitate H-sellers' strategy until $\bar{\kappa}$, posting $\{p^j\}_{j=\kappa+1}^{\bar{\kappa}}$, and θ_L in period $\bar{\kappa}+1$. Denote with σ_l' this imitating strategy. In equilibrium σ_l' cannot provide a strictly higher expected payoff than σ_L^*, i.e. for every $h_l^\kappa \in H^\kappa$:

$$\begin{aligned}
V_L^{\kappa+1}(\sigma^*,\pi^*|h_l^\kappa) \geq & \sum_{j=\kappa+1}^{\bar{\kappa}} \mathbb{P}_L^{\sigma_l'\sigma_{-l}^*\pi^*}\left(z(p^j)|\left(h_l^\kappa,\{p^s\times R\}_{s=\kappa+1}^{j}\right)\right)[p^j - v_L - (j-\kappa+1)c] \\
+ & \left[1-\sum_{j=\kappa+1}^{\bar{\kappa}} \mathbb{P}_L^{\sigma_l'\sigma_{-l}^*\pi^*}\left(z(p^j)|\left(h_l^\kappa,\{p^s\times R\}_{s=\kappa+1}^{j}\right)\right)\right][\theta_L - v_L - (\bar{\kappa}-\kappa+1)c] - c \\
\geq & \sum_{j=\kappa+1}^{\bar{\kappa}} \mathbb{P}_L^{\sigma_l'\sigma_{-l}^*\pi^*}\left(z(p^j)|\left(h_l^\kappa,\{p^s\times R\}_{s=\kappa+1}^{j}\right)\right)[v_H - v_L - (j-\kappa+1)c] \\
+ & \left[1-\sum_{j=\kappa+1}^{\bar{\kappa}} \mathbb{P}_L^{\sigma_l'\sigma_{-l}^*\pi^*}\left(z(p^j)|\left(h_l^\kappa,\{p^s\times R\}_{s=\kappa+1}^{j}\right)\right)\right][\theta_L - v_L - (\bar{\kappa}-\kappa+1)c] - c
\end{aligned} \tag{1.12}$$

Equations (1.11) and (1.12) imply:

$$(\bar{\kappa}-\kappa+1)c \geq \sum_{j=\kappa+1}^{\bar{\kappa}} \mathbb{P}_L^{\sigma_l'\sigma_{-l}^*\pi^*}\left(z(p^j)|\left(h_i^\kappa,\{p^s\times R\}_{s=\kappa+1}^{j}\right)\right)[v_H - \theta_L + (\bar{\kappa}-j)c] \tag{1.13}$$

As $\kappa < \bar{\kappa}$, starting from history $\tilde{h}_i^\kappa \in H^\kappa$ it holds $\sum_{j=\kappa+1}^{\bar{\kappa}} \mathbb{P}_H^{\sigma^*\pi^*}\left(z(p^j)|\left(\tilde{h}_i^\kappa,\{p^s\times R\}_{s=\kappa+1}^{j}\right)\right) > 0$. Hence, for every L-seller $i \in S_L^\kappa$ it must also hold:

$$\sum_{j=\kappa+1}^{\bar{\kappa}-1} \mathbb{P}_L^{\sigma_l'\sigma_{-l}^*\pi^*}\left(z(p^j)|\left(h_l^\kappa,\{p^s\times R\}_{s=\kappa+1}^{j}\right)\right)[v_H - \theta_L] > 0$$

because previous price history $h_l^{\kappa-1}$ is not observable to buyers and, once matched with a L-seller, they receive every signal $\xi \in \{H,L\}$ with positive probability. Therefore, inequality (1.13) cannot hold for c sufficiently small. ■

Lemma 1.9.2 *For c be sufficiently small, if there exists an equilibrium path in which H-sellers in $S_H^\kappa \neq \emptyset$ post a set of prices $P \subset \mathbb{R}_+$, all accepted with positive probability, then L-sellers in $S_L^\kappa \neq \emptyset$ post at least one price in P, unless they move to cohort $S^{\kappa+1}$ with probability one.*

Proof Lemma 1.9.2.

Consider sellers $i \in S_H^\kappa$ and $l \in S_L^\kappa$ with histories $h_i^{\kappa-1}$ and $h_l^{\kappa-1}$. Suppose on the contrary that L-sellers in S_L^κ do not post any price in P. Then, they can only trade at price θ_L. A seller $l \in S_L^\kappa$ prefers to post θ_L rather than the deviation strategy $\sigma_l'^\kappa(p|h_l^{\kappa-1}) = 1$ for some $p \in P$ only if:

$$\begin{aligned}\theta_L - v_L - c \geq\ & \mathbb{P}_L^{\sigma_l' \sigma_{-l}^* \pi^*}(z(p)|h_l^{\kappa-1})\,[p - v_L] + \left(1 - \mathbb{P}_L^{\sigma_l' \sigma_{-l}^* \pi^*}(z(p)|h_l^{\kappa-1})\right) V_L^{\kappa+1}(\sigma^*, \pi^*|h_l^{\kappa}) - c \\ \geq\ & \mathbb{P}_L^{\sigma_l' \sigma_{-l}^* \pi^*}(z(p)|h_l^{\kappa-1})\,[v_H - v_L] + \left(1 - \mathbb{P}_L^{\sigma_l' \sigma_{-l}^* \pi^*}(z(p)|h_l^{\kappa-1})\right) [\theta_L - v_L - c] - c\end{aligned}$$

The second inequality holds because L-sellers can always trade at θ_L in period $\kappa + 1$, and $p \geq v_H$ because H-sellers only trade at prices greater or equal to v_H. The inequality can be rewritten as:

$$c\left(1 - \mathbb{P}_L^{\sigma_l' \sigma_{-l}^* \pi^*}(z(p)|h_l^{\kappa-1})\right) \geq \mathbb{P}_L^{\sigma_l' \sigma_{-l}^* \pi^*}(z(p)|h_l^{\kappa-1})(v_H - \theta_L) \tag{1.14}$$

If p is accepted with positive probability, $\mathbb{P}_L^{\sigma_l' \sigma_{-l}^* \pi^*}(z(p)|h_l^{\kappa-1}) > 0$ as L-sellers may receive every signal $\xi \in \{H, L\}$, and buyers do not observe previous histories $h_i^{\kappa-1}$ and $h_l^{\kappa-1}$. By hypothesis, $\mathbb{P}_H^{\sigma^* \pi^*}(z(p)|h_i^{\kappa-1}) > 0$ and inequality (1.14) cannot hold for c sufficiently small. ■

Lemma 1.9.3 *For c sufficiently small, an equilibrium assessment (σ^*, π^*) satisfies:*

$$V_L^{\kappa}(\sigma^*, \pi^*|h_l^{\kappa-1}) - V_H^{\kappa}(\sigma^*, \pi^*|h_i^{\kappa-1}) \leq v_H - v_L \qquad \forall \kappa \in \mathbb{N}_0 \text{ and } \forall h_i^{\kappa-1} \in H^{\kappa-1}, \forall h_l^{\kappa-1} \in H^{\kappa-1}$$

The condition holds with equality only if: (i) H- and L-sellers post a common price and trade with probability one; or (ii) H- and L-sellers are both indifferent between postponing trade and posting a price p accepted only after a high signal.

Proof Lemma 1.9.3.
Consider an equilibrium (σ^*, π^*) and sellers $i \in S_H^{\kappa}$ and $l \in S_L^{\kappa}$. If seller $i \in S_H^{\kappa}$ prefers to stay out of the market it trivially holds:

$$V_H^{\kappa}(\sigma^*, \pi^*|h_i^{\kappa-1}) = 0 \qquad\qquad V_L^{\kappa}(\sigma^*, \pi^*|h_l^{\kappa-1}) = \theta_L - v_L - c$$

By Lemma 1.8.2 it holds for every seller in S_H^{κ} and S_L^{κ}, irrespective of previous histories.

Let:

$$\begin{aligned} K_H^N(\kappa) &:= \left\{ j \geq \kappa : \forall h_i^{j-1} \in H^{j-1} \text{ and } \forall p \text{ with } \sigma_H^{*^j}(p|h_i^{j-1}) > 0 \text{ it holds } \sigma_B^*(A|\mathscr{I}(p, \cdot, \xi)) = 0 \right\} \\ K_L^N(\kappa) &:= \left\{ j \geq \kappa : \forall h_l^{j-1} \in H^{j-1} \text{ and } \forall p \text{ with } \sigma_L^{*^j}(p|h_l^{j-1}) > 0 \text{ it holds } \sigma_B^*(A|\mathscr{I}(p, \cdot, \xi)) = 0 \right\} \end{aligned}$$

include all $j \geq \kappa$ such that all sellers in S_λ^j only post prices rejected with probability one. For simplicity, I denote with n the action to post a price rejected with probability one.

By definition, if $j \in K_\lambda^N(\kappa)$ for every $z \in Z^j(p)$ it is $\mathbb{P}_\lambda^{\sigma^*\pi^*}(z|h^{\kappa-1}) = 0$. By Lemma 1.9.1, it is $K_H^N(\kappa) \subseteq K_L^N(\kappa)$. By Lemma 1.9.2, if $j \notin K_H^N(\kappa)$ then either L-sellers do not trade with probability one, i.e. $j \in K_L^N(\kappa)$, or there exist $h_l^{j-1} \in H^{j-1}$, $h_i^{j-1} \in H^{j-1}$, and p^j such that $\sigma_L^{*j}(p^j|h_l^{j-1}) > 0$, $\sigma_H^{*j}(p^j|h_i^{j-1}) > 0$ and buyers accept p^j with positive probability. Consider such a sequence $\{p^j\}_{j\geq\kappa}$, $j \notin K_H^N(\kappa)$.

For $\{p^j\}_{j\geq\kappa}$, $j \notin K_H^N(\kappa)$, let $\bar{\kappa} = \arg\max_{j\in\mathbb{N}_0} \sigma_H^{*j}(p^j|h_i^{j-1}) > 0$, be the latest period of market participation for H-sellers along this price path. Denote with $\tilde{h}_i^{j-1}$ ($\tilde{h}_l^{j-1}$), $j \geq \kappa$, the history for seller $i \in S_H^\kappa$ ($l \in S_L^\kappa$) in which buyers reject price p^j if $j \notin K_H^N(\kappa)$ ($j \notin K_L^N(\kappa)$), or sellers play n if $j \in K_H^N(\kappa)$ ($K_L^N(\kappa)$). For $\kappa > \bar{\kappa}$, L-sellers' optimal best response is to post price θ_L and trade, so $\tilde{h}_L^{\bar{\kappa}+1} = (\tilde{h}_L^{\bar{\kappa}}, \theta_L \times A)$.

By Lemma 1.8.2, expected payoffs $V_\lambda^\kappa(\sigma^*,\pi^*)$ are independent from previous histories. Therefore, in equilibrium, sellers in S_λ^j are indifferent among all actions played with positive probability by any seller in S_λ^j. As a result, the price path along histories $\tilde{h}_i^{\bar{\kappa}}$ and $\tilde{h}_l^{\bar{\kappa}+1}$ provides an expected payoff equal to $V_H^\kappa(\sigma^*,\pi^*)$ and $V_L^\kappa(\sigma^*,\pi^*)$, respectively. Hence, it is possible to express sellers' expected payoff as:

$$
\begin{aligned}
V_H^\kappa(\sigma^*,\pi^*|h_i^{\kappa-1}) &= \sum_{j=\kappa}^{\bar{\kappa}} \mathbb{P}_H^{\sigma^*\pi^*}(z^j(p^j)|h_i^{\kappa-1})u_H^\kappa(z^j(p^j)) = \sum_{j\notin K_H^N(\kappa)}^{\bar{\kappa}} \mathbb{P}_H^{\sigma^*\pi^*}(z^j(p^j)|h_i^{\kappa-1})u_H^\kappa(z^j(p^j)) \\
V_L^\kappa(\sigma^*,\pi^*|h_l^{\kappa-1}) &= \sum_{j\notin K_L^N(\kappa)}^{\bar{\kappa}} \mathbb{P}_L^{\sigma^*\pi^*}(z^j(p^j)|h_l^{\kappa-1})u_L^\kappa(z^j(p^j)) + \mathbb{P}_L^{\sigma^*\pi^*}(z^{\bar{\kappa}+1}(\theta_L)|h_l^{\kappa-1})u_L^\kappa(z^{\bar{\kappa}+1}(\theta_L))
\end{aligned}
$$

To prove the statement I briefly state two preliminary observations:

a. $\mathbb{P}_L^{\sigma^*\pi^*}(z^{\bar{\kappa}+1}(\theta_L)|h_l^{\kappa-1}) \leq \left[1 - \sum_{j\notin K_L^N(\kappa)}^{\bar{\kappa}} \mathbb{P}_L^{\sigma^*\pi^*}(z^j(p^j)|h_l^{\kappa-1})\right]$

   Conditional on a history $h^{\kappa-1}$ the probability that a L-seller trades at price θ_L is *at least* equal to the complementary probability to trade in period j, $\kappa \leq j \leq \bar{\kappa}$, at price p^j along the equilibrium price path $\{p^j\}_{j=\kappa}^{\bar{\kappa}}$.

b. It holds:

$$
\sum_{j\notin K_L^N(\kappa)}^{\bar{\kappa}} \mathbb{P}_L^{\sigma^*\pi^*}(z^j(p^j)|h_l^{\kappa-1})u_H^\kappa(z^j(p^j)) - \sum_{j\notin K_H^N(\kappa)}^{\bar{\kappa}} \mathbb{P}_H^{\sigma^*\pi^*}(z^j(p^j)|h_i^{\kappa-1})u_H^\kappa(z^j(p^j)) \leq 0 \quad (1.15)
$$

   In equilibrium, a buyer accepts a posted price p^j either with probability one (for every ξ) or only if she receives a high signal $\xi = H$. Therefore, it is not possible to have:

$$
\sum_{j\notin K_L^N(\kappa)}^{\bar{\kappa}} \mathbb{P}_L^{\sigma^*\pi^*}(z^j(p^j)|h_l^{\kappa-1})u_H^\kappa(z^j(p^j)) > \sum_{j\notin K_H^N(\kappa)}^{\bar{\kappa}} \mathbb{P}_H^{\sigma^*\pi^*}(z^j(p^j)|h_i^{\kappa-1})u_H^\kappa(z^j(p^j)) = V_H^\kappa(\sigma^*,\pi^*|h_i^{\kappa-1})
$$

   If it were the case, a H-seller could profitably deviate by decreasing the probability of trading at earlier times in favour of later periods.[55]

Thanks to a. and b. it is possible to conclude that:

$$
\begin{aligned}
V_L^\kappa(\sigma^*,\pi^*|h_l^{\kappa-1}) \; - V_H^\kappa(\sigma^*,\pi^*|h_i^{\kappa-1}) &= \sum_{j\notin K_L^N(\kappa)}^{\bar\kappa} \mathbb{P}_L^{\sigma^*\pi^*}(z^j(p^j)|h_l^{\kappa-1})u_L^\kappa(z^j(p^j)) \\
&+ \mathbb{P}_L^{\sigma^*\pi^*}(z^{\bar\kappa+1}(\theta_L)|h_l^{\kappa-1})u_L^\kappa(z^{\bar\kappa+1}(\theta_L)) - \sum_{j\notin K_H^N(\kappa)}^{\bar\kappa} \mathbb{P}_H^{\sigma^*\pi^*}(z^j(p^j)|h_i^{\kappa-1})u_H^\kappa(z^j(p^j)) \\
&\leq \sum_{j\notin K_L^N(\kappa)}^{\bar\kappa} \mathbb{P}_L^{\sigma^*\pi^*}(z^j(p^j)|h_l^{\kappa-1})(v_H-v_L) + \left[1-\sum_{j\notin K_L^N(\kappa)}^{\bar\kappa} \mathbb{P}_L^{\sigma^*\pi^*}(z^j(p^j)|h_l^{\kappa-1})\right](\theta_L-v_L) \\
&+ \sum_{j\notin K_L^N(\kappa)}^{\bar\kappa} \mathbb{P}_L^{\sigma^*\pi^*}(z^j(p^j)|h_l^{\kappa-1})u_H^\kappa(z^j(p^j)) - \sum_{j\notin K_H^N(\kappa)}^{\bar\kappa} \mathbb{P}_H^{\sigma^*\pi^*}(z^j(p^j)|h_i^{\kappa-1})u_H^\kappa(z^j(p^j)) \\
&\leq v_H - v_L
\end{aligned}
$$

The first inequality holds because: (i) $\theta_L - v_L \geq u_L^\kappa(z^{\bar\kappa+1}(\theta_L))$;
(ii) $\mathbb{P}_L^{\sigma^*\pi^*}(z^{\bar\kappa+1}(\theta_L)|h_l^{\kappa-1}) \leq \left[1-\sum_{j\notin K_L^N(\kappa)}^{\bar\kappa} \mathbb{P}_L^{\sigma^*\pi^*}(z^j(p^j)|h_l^{\kappa-1})\right]$; and (iii) for every $z^j(p)\in Z^j$ it holds $u_L^\kappa(z^j(p)) = u_H^\kappa(z^j(p)) + v_H - v_L$. The second inequality follows from $\theta_L < v_H$ and equation (1.15).

Lastly, I show when it is $V_L^\kappa(\sigma^*,\pi^*|h_i^{\kappa-1}) - V_H^\kappa(\sigma^*,\pi^*|h_i^{\kappa-1}) = v_H - v_L$. For simplicity, I write V_λ^κ rather than $V_\lambda^\kappa(\sigma^*,\pi^*|h_i^{\kappa-1})$. Case (i) in the Lemma statement is immediate because $V_\lambda^\kappa = p - v_\lambda - c$. To prove case (ii) consider that, if a price p^κ is not accepted with probability one, then in equilibrium it is accepted only if $\xi = H$. The indifference condition requires:

$$
V_L^\kappa - V_H^\kappa = (1-\gamma)(p^\kappa - v_L) + \gamma V_L^{\kappa+1} - \gamma(p^\kappa - v_H) - (1-\gamma)V_H^{\kappa+1} = v_H - v_L
$$

Hence $p^\kappa = \frac{1}{2\gamma-1}\left[\gamma\left(V_L^{\kappa+1}+v_L\right) - (1-\gamma)\left(V_H^\kappa + v_H\right)\right]$.
Suppose, per contra, a type λ is not indifferent and strictly prefers to post p^κ. Then,

$$
\mathbb{P}_\lambda(H)(p^\kappa - v_\lambda) + (1-\mathbb{P}_\lambda(H))V_\lambda^{\kappa+1} - c > V_\lambda^{\kappa+1} - c
$$

i.e. $p^\kappa - v_\lambda > V_\lambda^{\kappa+1}$. Substituting p^κ into this expression and solving:

$$
\frac{\mathbb{P}_\lambda(H)}{2\gamma-1}\left[V_L^{\kappa+1} - V_H^{\kappa+1} - (v_H - v_L)\right] > 0
$$

which cannot hold as $V_L^{\kappa+1} - V_H^{\kappa+1} \leq v_H - v_L$. ■

Proof Lemma 1.4.1.
To prove the statement I show that there exists a $c^* > 0$ such that $\forall c \leq c^*$ no other be-

[55] For example, he could play an out-of-equilibrium strategy that includes prices rejected with probability one.

havioural strategy is admissible as an equilibrium of the game. By Proposition 1.4.2, no separating equilibrium exists for $c < \frac{1-\gamma}{\gamma}(\theta_H - \theta_L)$, otherwise L-sellers would deviate. By Lemma 1.8.2 I simplify notation and just write $V_\lambda^\kappa(\sigma^*, \pi^*)$ throughout the proof.

Step 1. If $q^\kappa < 1$ there is no equilibrium path in which H- and L-sellers in S^κ both use a mixed behavioural strategy.

Buyers play best responses in pure strategies and, for any posted price, they can either (i) accept with probability one; (ii) accept only if $\xi = j$, $j \in \{H, L\}$; or (iii) reject with probability one. Their best responses cannot depend on sellers' histories because they are not observable. Sequential rationality implies that—in equilibrium—no seller mixes between two different prices accepted with identical, positive, probability as he would strictly prefer the highest price. If seller $i \in S_\lambda^\kappa$ mixes among prices p_1, p_2, p_3, and buyers play $\sigma_B^\kappa(A|\mathscr{I}(p_1,\cdot,\xi)) = 1$ for every ξ, $\sigma_B(A|\mathscr{I}(p_2,\cdot,j)) = 1$ only for signal $\xi = j$ and reject otherwise, and $\sigma_B^\kappa(A|\mathscr{I}(p_3,\cdot,\xi) = 0$ for every ξ, then the following indifference conditions must hold:

$$\begin{aligned} p_1 - v_\lambda - c = &\ \mathbb{P}_\lambda(j)(p_2 - v_\lambda) + (1 - \mathbb{P}_\lambda(j)) V_\lambda^{\kappa+1}(\sigma^*, \pi^*) - c \\ p_1 - v_\lambda - c = &\ V_\lambda^{\kappa+1}(\sigma^*, \pi^*) - c \end{aligned}$$

However, this system of equations implies $p_1 = p_2$, but an identical price cannot be accepted with different probabilities. Therefore, in equilibrium, H- and L-sellers can only mix between: (i) two prices accepted with positive probability; (ii) one price accepted with positive probability and one (or more) rejected with probability one.

(i) Assume H-sellers mix between two prices (p_1, p_2). Buyers always accept p_2, but they accept p_1 only after signal $\xi = j, j \in \{H, L\}$. Mixing requires to be indifferent:

$$\mathbb{P}_H(j)(p_1 - v_H) + (1 - \mathbb{P}_H(j)) V_H^{\kappa+1}(\sigma^*, \pi^*) - c = p_2 - v_H - c \tag{1.16}$$

Moreover, H-sellers should prefer to trade rather than to move to period $\kappa + 1$, hence:

$$p_2 - v_H \geq V_H^{\kappa+1}(\sigma^*, \pi^*) \qquad p_1 - v_H \geq V_H^{\kappa+1}(\sigma^*, \pi^*) \tag{1.17}$$

In turn, inequalities (1.16) and (1.17) imply $p_1 \geq p_2 \geq v_H$. The inequality $p_1 \geq p_2$ holds strictly because buyers cannot accept the same price with different probabilities. Buyers accept only after signal j and reject otherwise only if $j = H$. Indeed, when H- and L-sellers mix on the same prices and $q^\kappa < 1$ buyers' beliefs satisfy $\pi_B(p_1,\cdot,H) > \pi_B(p_1,\cdot,L)$; as a result, $\mathbb{E}_{\pi_B}[\theta|\mathscr{I}(p_1,\cdot,H)] > \mathbb{E}_{\pi_B}[\theta|\mathscr{I}(p_1,\cdot,L)]$. If $j = L$ they would always accept p_1.

By hypothesis also L-sellers mix. Let's consider each possible mixed strategy.

a. L-sellers mix between two prices both accepted with positive probability. It is easy to realize that L-sellers pool on H-sellers' prices (p_1, p_2). Otherwise, in equilibrium, they can trade only at price θ_L; however, it would be a profitable deviation to trade with probability one posting $p_2 \geq v_H > \theta_L$. As H- and L-sellers mix on (p_1, p_2), the following indifference conditions hold:

$$V_H^{\kappa}(\sigma^*,\pi^*) = \quad \gamma(p_1 - v_H) + (1-\gamma)\, V_H^{\kappa+1}(\sigma^*,\pi^*) - c = p_2 - v_H - c$$
$$V_L^{\kappa}(\sigma^*,\pi^*) = \quad (1-\gamma)(p_1 - v_L) + \gamma V_L^{\kappa+1}(\sigma^*,\pi^*) - c = p_2 - v_L - c$$

As $p_1 > p_2$ each equation implies $V_\lambda^{\kappa+1}(\sigma^*,\pi^*) < p_2 - v_\lambda$ for $\lambda \in \{H, L\}$. Using the first equation, I get $p_1 = \frac{1}{\gamma}[p_2 - (1-\gamma)(v_H + V_H^{\kappa}(\sigma,\pi)]$. Substituting this expression into the second equation and simplifying:

$$p_2 = v_H + V_H^{\kappa+1}(\sigma,\pi) - \frac{\gamma^2}{2\gamma-1}\left[(v_H - v_L) - \left(V_L^{\kappa+1}(\sigma,\pi) - V_H^{\kappa+1}(\sigma,\pi)\right)\right]$$

By Lemma 1.9.3 $V_L^{\kappa+1}(\sigma,\pi) - V_H^{\kappa+1}(\sigma,\pi) \leq v_H - v_L$. Hence $V_H^{\kappa+1}(\sigma,\pi) \geq p_2 - v_H$ contradicting the previous implication $V_H^{\kappa+1}(\sigma,\pi) < p_2 - v_H$.

b. L-sellers mix between a price accepted with positive probability and a price rejected with probability one. As in point a. it is straightforward to realize L-sellers play either p_1 or p_2. Two cases are possible:

   1. L-sellers mix between p_2 and no trade. The indifference condition requires:
   $$p_2 - c = V_L^{\kappa+1} - c$$
   Moreover, posting p_1 is not a profitable deviation, hence:
   $$(1-\gamma)(p_1 - v_L) + \gamma V_L^{\kappa+1}(\sigma^*,\pi^*) - c \leq V_L^{\kappa+1}(\sigma^*,\pi^*) - c$$
   i.e. $p_1 - v_L \leq V_L^{\kappa+1}(\sigma^*,\pi^*)$, contradicting $p_1 > p_2$.

   2. L-sellers mix between p_1 and no trade. In turn, it must hold:
   $$(1-\gamma)(p_1 - v_L) + \gamma V_L^{\kappa+1}(\sigma^*,\pi^*) - c = V_L^{\kappa+1}(\sigma^*,\pi^*) - c$$
   i.e. $p_1 - v_L = V_L^{\kappa+1}(\sigma^*,\pi^*)$. In turn, using equation (1.17) and $p_1 > v_H$:
   $$V_L^{\kappa+1}(\sigma^*,\pi^*) - V_H^{\kappa+1}(\sigma^*,\pi^*) > v_H - v_L$$

contradicting Lemma 1.9.3.

(ii) H-sellers mix between a price accepted with positive probability and no trade. Therefore one of the two following indifference conditions holds:

$$\begin{aligned} \gamma(p_1 - v_H) + (1-\gamma)V_H^{\kappa+1}(\sigma^*, \pi^*) - c &= V_H^{\kappa+1}(\sigma^*, \pi^*) - c \\ p_2 - v_H - c &= V_H^{\kappa+1}(\sigma^*, \pi^*) - c \end{aligned}$$

Rearranging, I get $p_1 - v_H = V_H^{\kappa+1}(\sigma^*, \pi^*)$ and $p_2 - v_H = V_H^{\kappa+1}(\sigma^*, \pi^*)$, respectively. L-sellers' mixed strategy may either (a) offer two prices accepted with positive probability; or (b) offer a price accepted with positive probability or postpone trade.

a. Assume L-sellers mix between $(p_{1,L}, p_{2,L})$. Without loss of generality assume $p_{1,L}$ is accepted only after a signal $j \in \{H, L\}$ and $p_{2,L}$ is always accepted. Clearly, it is not optimal to post both prices different from H-sellers' one. Indeed, L-sellers would prefer to deviate and pool on H-sellers' price because it is accepted with the same probability of one between $p_{1,L}$ and $p_{2,L}$, but it provides a higher payoff as $\min\{p_1, p_2\} \geq v_H$. Therefore, L-sellers strategy can only mix between θ_L and the price posted by H-sellers. If H-sellers post p_2, it is never a best response to mix between θ_L and $p_2 \geq v_H > \theta_L$ as the latter is accepted with probability one. If H-sellers post p_1, L-sellers' indifference condition is:

$$\mathbb{P}_L(j)(p_1 - v_L) + (1 - \mathbb{P}_L(j))\, V_L^{\kappa+1}(\sigma^*, \pi^*) - c = \theta_L - v_L - c \tag{1.18}$$

Since $p_1 \geq v_H$, equation (1.18) implies $V_L^{\kappa+1}(\sigma^*, \pi^*) \leq \theta_L - v_L - \frac{\mathbb{P}_L(j)}{1-\mathbb{P}_L(j)}(v_H - \theta_L)$. L-sellers' continuation value always satisfies $V_L^{\kappa+1}(\sigma^*, \pi^*) \geq \theta_L - v_L - c$, because they can always trade at price θ_L. Therefore, for c sufficiently small both inequalities cannot contemporaneously hold.

b. Assume L-sellers mix between a price accepted with positive probability and no trade. As in point a., L-sellers pool on H-sellers' posted price. Hence, one of these two equations holds:

$$\begin{aligned} V_L^{\kappa}(\sigma^*, \pi^*) = (1-\gamma)(p_1 - v_L) + \gamma V_L^{\kappa+1}(\sigma^*, \pi^*) - c &= V_L^{\kappa+1}(\sigma^*, \pi^*) - c \\ V_L^{\kappa}(\sigma^*, \pi^*) = p_2 - v_L - c &= V_L^{\kappa+1}(\sigma^*, \pi^*) - c \end{aligned} \tag{1.19}$$

i.e. $p_1 - v_L = V_L^{\kappa+1}(\sigma^*, \pi^*)$ or $p_2 - v_L = V_L^{\kappa+1}(\sigma^*, \pi^*)$, respectively.
H-sellers indifference conditions imply $p_1 - v_H = V_H^{\kappa+1}(\sigma^*, \pi^*)$ or $p_2 - v_H = V_H^{\kappa+1}(\sigma^*, \pi^*)$. Hence, $V_L^{\kappa+1}(\sigma^*, \pi^*) - V_H^{\kappa+1}(\sigma^*, \pi^*) = v_H - v_L$, and

$V_L^\kappa(\sigma^*,\pi^*) - V_H^\kappa(\sigma^*,\pi^*) = v_H - v_L$. By Lemma 1.9.3, the only admissible strategies for sellers in S^j, $j > \kappa$, which satisfy this condition on continuation values are: (i) H- and L-sellers mix between a commonly posted price p^j and no trade; or (ii) H- and L-sellers in S^j post the same price p^j which is accepted with probability one. If sellers play strategy (i), $V_\lambda^j(\sigma^*,\pi^*)$ and p^j have to increase by c from j to $j+1$; see equation (1.19). Under TMN this strategy profile is not an equilibrium because sellers from different cohorts must have the same continuation value. Let's restrict attention to the TMO case. Observe that strategy (i) cannot be played for every $j \geq \kappa$, as it would eventually require to go above the upper bound $V_\lambda^j(\sigma^*,\pi^*) = \theta_H - v_\lambda - c$. Therefore, eventually sellers have to play strategy (ii); let $j^* \geq \kappa$ denote this future period. In turn, this requires to have $q^{j^*} > q_c^{IP}$. Clearly, the price posted under strategy (i) for $j < j^*$ has to be accepted only if $\xi = H$.

By assumption, $q^0 < q_c^S < q_c^{IP}$. In order to increase H-sellers' share from q^0 to $q^{j^*} \geq q_c^{IP}$, sellers should play a behavioural strategy of type (i) for $j^* - 1$ initial periods in order to increase q^j from q^0 to at least q_c^{IP}; then all sellers in S^{j^*} should trade with probability one at p^{j^*}. Let α^j and β^j denote the probability that H- and L-sellers, respectively, play p^j for $j < j^*$; the complementary probability denotes the probability to post a price rejected with probability one. In equilibrium, for $j < j^*$ buyers accept p^j after $\xi = H$, i.e.:

$$\frac{q^j\alpha^j\gamma\theta_H + (1-q^j)\beta^j(1-\gamma)\theta_L}{q^j\alpha^j\gamma + (1-q^j)\beta^j(1-\gamma)} \geq p^j \;\rightarrow\; \frac{\beta^j}{\alpha^j} \leq \frac{q^j}{1-q^j}\frac{\gamma}{1-\gamma}\frac{\theta_H - p^j}{p^j - \theta_L} \tag{1.20}$$

The share q^j increases only if:

$$\frac{q^{j+1}}{1-q^{j+1}} = \frac{q^j[(1-\alpha^j) + \alpha^j(1-\gamma)]}{(1-q^j)[(1-\beta^j) + \beta^j\gamma]} > \frac{q^j}{1-q^j} \quad \rightarrow \quad \frac{\beta^j}{\alpha^j} > \frac{\gamma}{1-\gamma} \tag{1.21}$$

Equations (1.21) and (1.20) together imply:

$$q^j\theta_H + (1-q^j)\theta_L > p^j \geq v_H$$

This is possible only if $q^0 > \frac{v_H - \theta_L}{\theta_H - \theta_L}$. However, by assumption $q^0 < q_c^S = \frac{v_H - \theta_L + c}{\theta_H - \theta_L}$ and for c small enough the two inequalities cannot both hold.

Step 2. There is no equilibrium path in which L-sellers in $S_L^\kappa \neq \emptyset$ only post prices rejected with probability one, and H-sellers in $S_H^\kappa \neq \emptyset$ trade with positive probability.[56]

[56] Notice that it is a more restrictive statement than Lemma 1.9.2.

H-sellers in S_H^κ post a price p accepted with positive probability, so $p \leq \mathbb{E}_{\pi_B}[\theta|\mathscr{I}(p,\cdot,\xi)] = \theta_H$ for all $\xi \in \{H,L\}$, because, by hypothesis, L-sellers do not pool on this price. If $p < \theta_H$ all buyers accept with probability one. H-sellers' possible behavioural strategies are:

a. H-sellers post $p \leq \theta_H$ with probability one and buyers always accept. Therefore, it must hold $p - v_L - c \leq V_L^{\kappa+1} - c$ and $p - v_H - c \geq V_H^{\kappa+1} - c$. Both inequalities imply $V_L^{\kappa+1} - V_H^{\kappa+1} \geq v_H - v_L$. By Lemma 1.9.3 the inequality cannot hold strictly; moreover, the argument in Step 1(ii) b. excludes the equality case under TMO (when $q^0 < q_c^N$) and under TMN (always).

b. H-sellers post θ_H with probability one and buyers accept only for $\xi = j$, $j \in \{H,L\}$. L-sellers prefer to postpone trade only if:

$$V_L^{\kappa+1}(\sigma^*,\pi^*) \geq \mathbb{P}_L(j)(\theta_H - v_L) + (1 - \mathbb{P}_L(j))\, V_L^{\kappa+1}(\sigma^*,\pi^*)$$

i.e. $V_L^{\kappa+1}(\sigma^*,\pi^*) \geq \theta_H - v_L$. However, in every equilibrium an upper bound on the expected payoff is $V_L^{\kappa+1} \leq \theta_H - v_L - c$. A similar argument applies if L-sellers mix between two prices both accepted with positive probability.

c. H-sellers mix between a price p accepted with positive probability and no trade. If p is always accepted then H-sellers' indifference requires

$$p - v_H = V_H^{\kappa+1}(\sigma^*,\pi^*)$$

L-sellers do not trade at p if $V_L^{\kappa+1}(\sigma^*,\pi^*) \geq p - v_L$. Then,

$$V_L^{\kappa+1}(\sigma^*,\pi^*) - V_H^{\kappa+1}(\sigma^*,\pi^*) \geq v_H - v_L$$

If the inequality is strict then it is inconsistent with Lemma 1.9.3. If it holds with equality, see Step 1(ii) b.

If H-sellers mix between θ_H and no trade then:

$$\mathbb{P}_H(j)(\theta_H - v_H) + (1 - \mathbb{P}_H(j))\, V_H^{\kappa+1}(\sigma^*,\pi^*) = V_H^{\kappa+1}(\sigma^*,\pi^*)$$

i.e. $V_H^{\kappa+1}(\sigma^*,\pi^*) = \theta_H - v_H$. But in every equilibrium an upper bound on expected payoffs is $V_H^{\kappa+1}(\sigma^*,\pi^*) \leq \theta_H - v_H - c$.

Step 3. There is no equilibrium path in which sellers in $S_H^\kappa \neq \emptyset$ and $S_L^\kappa \neq \emptyset$ play semi-separating behavioural strategies, and all posted prices are accepted with positive probability.

Suppose, on the contrary, that such a behavioural strategy is played in equilibrium. Let's consider the two different cases:

(i) L-sellers mix and H-sellers play a pure strategy. The only relevant case to consider is when L-sellers mix between θ_L (accepted with probability one) and p (accepted only if $\xi = H$), and H-sellers only post p. L-sellers mix between the two prices only if:

$$V_L^{\kappa}(\sigma^*,\pi^*) = (1-\gamma)(p-v_L) + \gamma V_L^{\kappa+1}(\sigma^*,\pi^*) - c = \theta_L - v_L - c$$

L-sellers can always trade at θ_L so $V_L^{\kappa+1}(\sigma^*,\pi^*) \geq \theta_L - v_L - c$. Hence,

$$\theta_L - v_L - c \geq (1-\gamma)(p-v_L) + \gamma(\theta_L - v_L - c)$$

i.e. $\theta_L \geq p - \frac{c}{1-\gamma}$. For c small enough $p < v_H$ as $v_H > \theta_L$. H-sellers would not post this price in equilibrium as lower than their reservation value.

(ii) H-sellers mix and L-sellers play a pure strategy. Sellers in S^{κ} play this behavioural strategy only if H-sellers mix between θ_H (played with probability α and accepted only if $\xi = H$) and a price p^{κ} (always accepted). L-sellers post p^{κ} with probability one.

I provide a separate proof for each assumption on time on market observability.

a. Under TMO, buyers accept p^{κ} irrespective of signal ξ only if $\mathbb{E}_{\pi_B}[\theta|(p^{\kappa},\kappa,L)] \geq p^{\kappa}$. In turn, a necessary condition is $q^{\kappa} > q_c^{IP}$ (for $\alpha = 0$, it is higher for $\alpha > 0$).[57] As $q^0 < q_0^{IP} < q_c^{IP}$ it is sufficient to prove that $q^{\kappa} < q^0$ for every $\kappa \geq 1$. Consider $j \in \mathbb{N}_0$ such that $q^j < q_0^{IP}$. Therefore, in period j the proposed behavioural strategy cannot be played. By steps 1, 2 , 3 (i) and Proposition 1.4.2: (i) either both H- and L-sellers do not trade with probability one; or (ii) $\sigma_H^j(p) = \sigma_L^j(p) = 1$, $\sigma_B(A|\mathscr{I}_B(p,\kappa,H)) = 1$ and $\sigma_B(A|\mathscr{I}_B(p,\kappa,L)) = 0$. In case (i) $q^{j+1} = q^j$. In case (ii) buyers accept only if $\xi = H$, so H- and L-sellers trade with probability γ and $1-\gamma$, respectively. By the law of large numbers, a higher share of H-sellers trades and exits the market, so $q^{j+1} < q^j$. Let $j = 0$ concludes the argument.

b. Under TMN, buyers do not distinguish sellers in different cohorts S^{κ}. In equilibrium, two cohorts of sellers $S_{\lambda}^{\kappa'}$ and $S_{\lambda}^{\kappa''}$, $\kappa' \neq \kappa''$ cannot play strategy profiles leading to different expected payoffs. If this were the case, there would be a profitable deviation for one cohort of sellers as buyers cannot

[57] See section 1.4 for a characterization of q_c^{IP}. The subscript c refers to the associated search cost c.

observed previous prices. For every $\kappa \in \mathbb{N}_0$ such that $S^{\kappa}_{\lambda} \neq \emptyset$, it must hold $V^{\kappa}_{\lambda}(\sigma^*, \pi^*) = V_{\lambda}(\sigma^*, \pi^*)$, and H-sellers' indifference condition requires:

$$V_H(\sigma^*, \pi^*) = \gamma(\theta_H - v_H) + (1-\gamma)V_H(\sigma^*, \pi^*) - c = p^{\kappa} - v_H - c$$

hence $p^{\kappa} = \bar{p} = \theta_H - \frac{1-\gamma}{\gamma}c$.

Let $\bar{q} = \sum_{\kappa \in \mathbb{N}_0} \frac{S^{\kappa}}{S} q^{\kappa}$. Buyers accept $\bar{p}$ only if:

$$\frac{\bar{q}(1-\alpha)(1-\gamma)\theta_H + (1-\bar{q})\gamma\theta_L}{\bar{q}(1-\alpha)(1-\gamma) + (1-\bar{q})\gamma} \geq \theta_H - \frac{1-\gamma}{\gamma}c$$

hence,

$$\frac{\bar{q}}{1-\bar{q}} \geq \frac{\gamma}{(1-\gamma)(1-\alpha)} \frac{\theta_H - \theta_L - c}{c} \tag{1.22}$$

By Step 1, 2, 3(i), the only admissible behavioural strategy for sellers in cohorts S^j, $j \neq \kappa$, are: (i) both H- and L-sellers do not trade with probability one; (ii) H- and L-sellers post the same price p_H, say, and buyers accept only if $\xi = H$; (iii) H- and L-sellers post the same price p_L, say, and buyers accept with probability one; (iv) H-sellers mix between two prices and L-sellers post one price with probability one. Since all cohorts receive the same expected payoff, it must be $p_H = \theta_H$ and $p_L = \bar{p} = \theta_H - \frac{1-\gamma}{\gamma}c$. Importantly, the behavioural strategies (i), (ii) and (iii) imply that H-sellers in S^j_H trade at least with the same probability as L-sellers in S^j_L. As a result, $\bar{q}$ cannot be higher than q^0 if sellers only play strategies (i), (ii) and (iii). Therefore, if it is not possible to satisfy equation (1.22) in an equilibrium in which, for every $\kappa \in \mathbb{N}_0$, H-sellers in S^{κ}_H mix between θ_H and $\bar{p}$, and L-sellers in S^{κ}_L play $\bar{p}$ with probability one, then it is never possible. In the proposed equilibrium, stationarity requires exit and entry flows for each type of sellers to be equal. Denote with $\bar{S}$ the equilibrium mass of sellers:

$$\begin{cases} q^0 = \bar{S}\bar{q}[\alpha\gamma + (1-\alpha)] \\ (1-q^0) = \bar{S}(1-\bar{q}) \end{cases} \Rightarrow \quad \begin{cases} \bar{S} = \frac{1-q^0}{(1-\bar{q})} \\ \frac{\bar{q}}{1-\bar{q}} = \frac{1}{\alpha\gamma - (1-\alpha)} \frac{q^0}{1-q^0} \end{cases}$$

Therefore, this equilibrium exists only if equation (1.22) is satisfied, i.e.

$$\frac{1}{\alpha\gamma + (1-\alpha)} \frac{q^0}{1-q^0} \geq \frac{\gamma}{(1-\gamma)(1-\alpha)} \frac{\theta_H - \theta_L - c}{c}$$

This expression is more likely to hold for α close to zero. Therefore, a necessary condition is:

$$\frac{q^0}{1-q^0} \geq \frac{\gamma^2}{1-\gamma}\frac{\theta_H-\theta_L-c}{c} \tag{1.23}$$

For $q^0 = q_c^{IP} = \frac{\gamma}{1-\gamma}\frac{v_H+c-\theta_L}{\theta_H-v_H-c}$[58] equation (1.23) becomes:

$$\frac{v_H+c-\theta_L}{\theta_H-v_H-c} \geq \gamma\frac{\theta_H-\theta_L-c}{c}$$

However, for c sufficiently small this inequality cannot hold.

Step 4. There is no equilibrium path in which H- and L-sellers in $S_H^\kappa \neq \emptyset$ and $S_L^\kappa \neq \emptyset$ post a price accepted with probability one.

I provide a separate proof for each assumption on time on market observability.

a. Under TMO the argument is analogous to the one in Step 3 part (ii) point a.

b. Under TMN, an analogous argument to the one in Step 3 part (ii) point b establishes that all cohorts get the same continuation value. By Step 1, 2, 3 the only admissible behavioural strategy profiles for sellers in cohorts S^j, $j \neq \kappa$, are: (i) both H- and L-sellers do not trade with probability one; or (ii) H- and L-sellers post the same price p_H and buyers accept only if $\xi = H$; or (iii) H- and L-sellers post the same price p_L and buyers accept with probability one.

   By definition of q_c^{IP}, the behavioural strategy (iii) is possible only if $\bar{q} = \sum_{\kappa \in \mathbb{N}_0} \frac{S^\kappa}{S} q^\kappa \geq q_c^{IP}$, i.e. there must be at least one cohort $S^{\tilde{\kappa}}$, $\tilde{\kappa} \neq \kappa$, such that L-sellers trade with a higher probability than H-sellers. However, all admissible behavioural strategies (i), (ii) and (iii) imply that H-sellers' probability to trade is at least equal to L-sellers' one.

Step 5. If $S_H^\kappa \neq \emptyset$ then all sellers in S^κ post the same price. If time on market is not observable all sellers post the same price.

By steps 1, 2, 3, 4 and Lemma 1.9.2 the only admissible behavioural strategy profiles are: (i) H- and L-sellers only post prices rejected with probability one; (ii) H- and L-sellers in S^κ post the same price p^κ and buyers accept only if $\xi = H$. Under TMN, buyers do not observe sellers' cohort, and sellers cannot post different prices accepted with the same probability; hence, $p^\kappa = \bar{p}$, for all $\kappa \in \mathbb{N}_0$. Postponing trade does not change the future trade price but it increases search costs. As a result, they find strictly convenient to post $\bar{p}$ until they trade. ■

[58] See section 1.4 for a defnition.

Bibliography

AKERLOF, G. (1970): "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism," *Quarterly Journal of Economics Studies*, 84, 488–500.

ATAKAN, A. (2006): "Assortative Matching with Explicit Search Costs," *Econometrica*, 74, 667–680.

BAYE, M. R., MORGAN, J. AND P. SCHOLTEN (2004): "Price Dispersion in the Small and in the Large: Evidence from an Internet Price Comparison Site," *Journal of Industrial Economics*, 52, 463–496.

BLOUIN, M. (2003): "Equilibrium in a Decentralized Market with Adverse Selection," *Economic Theory*, 22, 245–62.

BLOUIN, M., AND R. SERRANO (2001): "A Decentralized Market with Common Values Uncertainty: Non-Steady States," *Review of Economic Studies*, 68, 323–346.

CAMARGO, B., AND B. LESTER (2014): "Trading Dynamics in Decentralized Markets with Adverse Selection," *Journal of Economic Theory*, 153, 534–568.

CLAY, K., KRISHNAN, R. AND E. WOLFF (2001): "Prices and Price Dispersion on the Web: Evidence from the Online Book Industry," *Journal of Industrial Economics*, 49, 521–539.

DALEY, B. AND B. GREEN (2012): "Waiting for News in the Market for Lemons," *Econometrica*, 80, 1433–1504.

DELLAROCAS, C. (2005): "Reputation Mechanisms Design in Online Trading Environments with Pure Moral Hazard,"*Information Systems Research*, 16, 209–230.

DEWAN, S., AND V. HSU (2004): "Adverse Selection in Electronic Markets: Evidence from Online Stamp Auctions,"*Journal of Industrial Economics*, 52, 497–516.

DIAMOND, P. (1971): "A Model of Price Adjustment," *Journal of Economic Theory*, 3, 156–68.

DUFFIE, D., N. GÂRLEANU, AND L.H. PEDERSEN (2005): "Over-the-Counter Markets," *Econometrica*, 73, 1815–1847.

DUFFIE, D., N. GÂRLEANU, AND L.H. PEDERSEN (2007): "Valuation in Over-the-Counter Markets," *Review of Finanical Studies*, 20, 1865–1900.

DUFFIE, D., MALAMUD, S., AND G. MANSO (2009): "Information Percolation With Equilibrium Search Dynamics," *Econometrica*, 77, 1513–1574.

DUFFIE, D., MALAMUD, S., AND G. MANSO (2014): "Information Percolation in Segmented Markets," *Journal of Economic Theory*, 153, 1–32.

FUDENBERG, D., AND J. TIROLE (1991): "Perfect Bayesian Equilibrium and Sequential Equilibrium," *Journal of Economic Theory*, 53, 236–260.

FUCHS, W., AND A. SKRZYPACZ (2013): "Cost and Benefits of Dynamic Trading in a Lemons Market," working paper.

FUCHS, W., AND A. SKRZYPACZ (2014): "Government Interventions in a Dynamic Market with Adverse Selection," working paper.

FUCHS, W., A. ÖRY AND A. SKRZYPACZ (2014): "Transparency and Distressed Sales under Asymmetric Information," working paper.

GALE, D. (1986a): "Bargaining and competition part I: Characterization," *Econometrica*, 54, 785–806.

GALE, D. (1986b): "Bargaining and competition part II: Existence," *Econometrica*, 54, 807–818.

GOLDMANIS, M., HORTACSU, A., SYVERSON, C. AND E. ÖNSEL (2010): "E-Commerce and the Market Structure of Retail Industries," *The Economic Journal*, 120, 651–682.

GUERRIERI, V., AND R. SHIMER (2014): "Dynamic Adverse Selection: A Theory of Illiquidity, Fire Sales and Flight to Quality," *American Economic Review*, 104, 1875–1908.

HENDEL, I., NEVO, A., AND F. ORTALO-MAGNÉ (2009): "The Relative Performance of Real Estate Marketing Platforms: MLS versus FSBOMadison.com.," *American Economic Review*, 99, 1878–98.

HENDON, E., JACOBSEN, H.J., AND B. SLOTH (1996): "The One-Shot Deviation Principle for Sequential Rationality," *Games and Economic Behaviour*, 12, 274–282.

HÖRNER, J., AND N. VIEILLE (2009): "Public vs. private offers in the market for lemons," *Econometrica*, 51, 29–69.

JANSSEN, M., AND S. ROY (2002): "Dynamic Trading in a Durable Good Market with Asymmetric Information," *International Economic Review*, 43, 257–82.

JANSSEN, M., AND S. ROY (2004): "On Durable Goods Markets with Entry and Adverse Selection," *Canadian Journal of Economics*, 37, 552–589.

JIN, G.Z., AND A. KATO (2007): "Dividing Online and Offline: A Case Study," *Review of Economic Studies*, 74, 981–1004.

JULLIEN, B., AND T. MARIOTTI (2006): "Auction and the Informed Seller Problem," *Games and Economic Behaviour*, 56, 225–258.

KAYA, A., AND K. KIM (2014): "Trading Dynamics in the Market for Lemons," working paper.

KIM, K. (2014): "Information about Sellers' Past Behaviour in the Market for Lemons," working paper.

LAGOS, R., AND G. ROCHETEAU (2009): "Liquidity in Asset Markets with Search Frictions," *Econometrica*, 77, 403–426.

LAUERMANN, S., AND A. WOLINSKY (2013): "Search with Adverse Selection," working paper.

LEWIS, G. (2011): "Asymmetric Information, Adverse Selection and Online Disclosure: The Case of eBay Motors," *American Economic Review*, 101, 1535–1546.

MAILATH, G., OKUNO-FUJIWARA, M., AND A. POSTLEWAITE (1993): "On Belief Based Refinements in Signalling Games," *Journal of Economic Theory*, 60, 241–276.

MILGROM, P., AND C. SHANNON (1994): "Monotone Comparative Statics," *Econometrica*, 62, 157–180.

MORENO, D., AND J. WOODERS (2010): "Decentralized Trade Mitigates the Lemons Problem," *International Economic Review*, 51, 383–399.

MORENO, D., AND J. WOODERS (2014): "Dynamic Market for Lemons: Performance, Liquidity, and Policy Intervention," working paper.

PHILIPPON, T., AND V. SKRETA (2012): "Optimal Interventions in markets with Adverse Selection," *American Economic Review*, 102, 1–28.

RUBINSTEIN, A., AND A. WOLINSKY (1985): "Equilibrium in a Market with Sequential Bargaining," *Econometrica*, 53, 1133–50.

SERRANO, R., AND O. YOSHA (1993): "Information Revelation in a Market with Pairwise Meetings: the One Sided Information Case," *Economic Theory*, 3, 481–499.

SPENCE, M. (1973): "Job Market Signalling," *The Quarterly Journal of Economics*, 87, 355–374.

TAYLOR, C. (1999): "Time-on-the-Market as a Sign of Quality," *Review of Economic Studies*, 66, 555–578.

TIROLE, J. (2012): "Overcoming Adverse Selection: How Public Intervention Can Restore Market Functioning," *American Economic Review*, 102, 19–59.

TUCKER, C., ZHANG, J., AND T. ZHU (2013): "Days on Market and Home Sales," *RAND Journal of Economics*, 44, 337–360.

VAYANOS, D., AND O. WEILL (2008): "A Search-based Theory of the On-the-run Phenomenon," *Journal of Finance*, 63, 1351–1389.

WOLINSKY, A. (1990): "Information Revelation in a Market with Pairwise Meetings," *Econometrica*, 58, 1–23.

Chapter 2

Peer Monitoring Incentives via Central Clearing Counterparties

Central Clearing Counterparties (CCPs) envisage that, upon a member's default, surviving members have to partially cover losses. I study how the way to distribute losses among CCP members change their incentives to peer monitor each other's. The optimal design exploits dealers' superior information on the credit risk of other CCP members. My results suggest that a higher share of losses should be paid by surviving members with a greater trade exposure to the defaulting dealer. In practice, this mechanism can be implemented through variation margin haircutting, or higher rights of assessment. If a CCP distributes losses without reference to previous trade, equilibrium outcomes may be inferior to what can be achieved with no clearing.

2.1 Introduction

The recent financial crisis brought to public attention the importance of over-the-counter (OTC) markets, and it spurred a global debate on the existing financial architecture. After Lehman Brothers' bankruptcy, OTC markets were often considered a source of systemic risk because of their lack of transparency, the inextricable network relationships, and the complexity of financial products exchanged. To enhance financial stability, the Dodd-Frank Act (2010) in the US, and the EMIR regulation (2013) in the EU, introduced several new legislative provisions on trade transparency and risk mitigation.

Central Clearing Counterparties (CCPs) are going to play an important role in the new financial architecture. Several OTC products will be subject to mandatory clearing through a CCP. Although CCPs were first introduced over a century ago, their role has been confined to the ordered execution and clearing of securities traded on regulated exchanges. Their novel use in a new range of financial instruments, such as interest rate swaps, credit derivatives, and repos, raises a novel set of questions on their effectiveness as an instrument to build a more resilient financial system.

A CCP performs several functions: multi-lateral netting, post-trade transparency, setting initial and variation margins, and loss mutualization in case of a member's default. Despite their relevance for financial stability, netting and trade transparency may be independently achieved through trade compression services and central trade repositories, respectively. In this paper, I am going to focus exclusively on one *distinctive* feature of CCPs: the design of the loss mutualization scheme. Indeed, a CCP can—upon the default of a member—distribute the resulting losses among its surviving members. In the current debate, this possibility has been considered relevant only to the extent it provides an additional safeguard for third-parties and other CCP members.

This paper looks at the loss allocation rules from a different perspective. Specifically, I consider how the design of the default waterfall[1] affects peer monitoring incentives among financial dealers. Similarly to Stiglitz (1990) and Varian (1990), I use the term 'peer monitoring' as the possibility of alleviating borrowers' moral hazard problem through peers' interaction. Peers usually have superior information relative to outsiders, and the goal is to design mechanisms that exploit this information to the benefit of outsiders.[2] In the context of my model, the default waterfall change dealers' incentives to trade with a risky counter-party, and, in turn, this affects dealers' agency

[1] Initial margins and default funds contributions are the most important tools of a broader system of safeguards named default waterfall (see Elliott (2013) for a discussion). This is a set of rules determining the hierarchy of funds used to cover losses from a member default, and the extent to which each surviving member is responsible for these losses.

[2] I assume dealers do not pay any cost to acquire superior information on their peers. Endogenous information acquisition is a valuable extension of the model and I plan to pursue it in the future.

problem with outside investors. I abstract from the risk insurance role of the default waterfall, and I focus exclusively on the interaction between its design and dealers' incentives to choose a risky business conduct. To stress this perspective, I sometimes refer to default fund contributions with the term *penalties*, to highlight their punishment role for surviving CCP members. My narrow interpretation of a CCP—a financial institution that sets collateral margins and imposes penalties contingent on default events—is functional to the specific research questions. Throughout the paper, I consider an inter-dealer market in which dealers know the identity of their counter-parties. This is often the case for OTC markets, while in trade platforms with anonymous clearing—such as the stock market—peer monitoring is not an option.

I address the following questions: Why do CCPs affect peer monitoring? Does their current design improve or discourage risky behaviour? What is the optimal default waterfall to maximize peer monitoring inventives? Does it depend on the correlation among dealers' shocks?

My analysis builds on two main considerations. First, a dealer finds essential to trade in the inter-dealer market, either to hedge his exposure with end users, or to implement his own proprietary investments.[3] The importance of inter-dealer markets is hardly disputable, as mirrored by their volume of transactions. Second, financial institutions—such as dealers—may have superior skills to assess the credit risk of a similar institution. They have a comprehensive understanding of the industry, real time information on market conditions, and possibly superior information on the risk exposure of other dealers, as they are often on the other side of a financial transaction. Although it is difficult to assess whether financial institutions have superior information on their peers, it is a widely shared view among policymakers, and it has been recently corroborated by some papers.[4]

I argue that the design of the default waterfall may significantly change peer monitoring incentives. To maximize their effectiveness—and decrease initial margins—a CCP should impose penalties *only* on the surviving dealers who had previously traded with a defaulting member. For OTC interest rate or foreign exchange swaps, this could be implemented with a variation margin haircut; by definition, a CCP runs a matched book: if a member loses, another member gains. For repos or credit derivatives, it could be implemented as a haircut on the principal to be reimbursed. However, only a few CCPs have recently introduced this loss allocation rule. All CCPs envisage to first

[3]Dealers aim to have a 'neutral' stance by systematically hedging their exposure on the inter-dealer market. They hold extremely large notional positions, even if their net exposure is much smaller. Despite a recent trend towards more open participation, this market is still very concentrated: a few dealers intermediate the vast majority of derivative contracts and repos. The OCC reports that in Q2 2013 the top four dealers (JP Morgan Chase, Bank of America, Citigroup and Goldman Sachs) held 93% of all derivative contracts in the US. Corporate end-users, retail banks, mutual and pension funds participate in the market mainly to take a directional position with respect to a specific risk.

[4]See Affinito (2012) and Bräuning et al. (2014).

exhaust the resources of a pre-funded default fund (or *guarantee fund*), and later to call members to contribute with additional funds (*rights of assessment*). These contributions are normally computed in approximate relation to the amount of risk that each dealer brings to the CCP. In general, the allocation of losses does not depend directly on dealers' previous trade intensity with the defaulting member. I show why this common default waterfall structure may reduce peer monitoring incentives, and increase other forms of costly disciplining devices such as initial margins.

I set up a simple risk-shifting model, and I extend it with an inter-dealer market. Dealers have perfect information on the default probabilities of other dealers. I characterize the equilibrium contract, and I evaluate how the introduction of a CCP affects market outcomes relative to an unregulated market. In my model, dealers face a standard risk-shifting problem: after signing an initial contract with lenders, they may prefer to undertake risky activities. A safe business conduct requires to pay an effort cost; it can be interpreted as the cost of implementing good risk management techniques, or to select good proprietary investments. Dealers may 'commit' to exert effort by posting initial collateral to investors, or, importantly, incentives may result from the interaction on the inter-dealer market. After the effort choice, a dealer's credit risk realizes but it is not verifiable. Outside investors and dealers differ in their ability to observe credit risk: investors have no information, while dealers observe their own and other dealers' default probabilities. Once credit risk information is observed, dealers look for a hedge. If two dealers agree to hedge, they both avoid to pay a cost borne when holding an unhedged position; however, they are also exposed to an exogenous risk of contagion when the hedging partner happens to default. For tractability reasons, I restrict attention to bilateral hedging, and I do not allow collateral posting in the inter-dealer market. Both extensions are interesting future research directions. I first solve the model with only two dealers, and I later extend my results to the multiple dealers case to highlight the new economic forces at work.

In my model, the inter-dealer market provides disciplining incentives through two economic channels. First, a risky dealer may not hedge—suffering a loss—if all other dealers prefer not to be exposed to a high risk of contagion.[5] The possibility *not* to be able to hedge provides an *ex ante* incentive to avoid risky investment activities. However, if it is too costly not to hedge, this threat is less credible. Even if the other dealer is perceived as risky, it may not be easy to find another hedge and, despite contagion risk, it may still be convenient to accept. In section 2.5, I point out a second disciplining mechanism, based on the endogenous peer selection. For two safe dealers, it is

[5]The default of a dealer poses a serious threat to his counter-parties, as a large exposure to a defaulter may trigger a contagion spiral. For some empirical evidence on contagion risk see Liedorp et al. (2010) and Craig et at. (2014).

mutually beneficial to hedge and avoid contagion risk. Therefore, risky dealers either match with a safe counter-party, risking not to hedge, or they hedge with another risky counter-party. However, matching with a risky dealer is costly because of the higher risk of contagion. For this reason, it is convenient to exercise effort—maximizing the chances of being safe—and increase the probability to match with another safe counter-party.

I also analyze when each peer monitoring mechanism is more effective. The threat of being excluded from the inter-dealer market is more relevant when there is a small number of dealers, and credit risk shocks among dealers are strongly correlated (for example, due to a macroeconomic deterioration). In contrast, the incentive to be safe, and hedge with another safe dealer, is more relevant when the inter-dealer market has a large number of participants, and there is a low correlation among dealers' credit risk shocks. In this case, it is likely that some dealers in the market are risky, and each dealer assigns a low probability to the event of matching with a dealer of different credit risk.

Both disciplining mechanisms work because hedging with a risky dealer leads to a higher risk of contagion, i.e. a dealer may lose his profits if the other counter-party defaults. A CCP can strengthen *both* disciplining mechanisms through the *same* design of the loss allocation rules. Upon the default of a clearing member, surviving dealers are called to cover losses. Transfers to the CCP increase the cost of matching with a risky dealer, as they impose losses even when no contagion takes place. Importantly, my results stress that penalties should be paid only if a dealer has previously traded with the defaulting member.[6] If a pair of dealers does not hedge, and one dealer defaults, no CCP member should be responsible to cover his losses. Indeed, if penalties were always paid—irrespective of the hedging decision—the threat of being excluded from the market would be less credible; as a result, the risk-shifting problem worsens, increasing the need to restore incentives through initial collateral.

As in other risk-shifting models, higher collateral requirements increase dealers' skin in the game, and alleviate the risk-shifting problem. However, this ex ante risk mitigation device is expensive,[7] and ex post penalties provide a less costly alternative. In my model the relationship between initial margins and peer monitoring incentives is not of pure substitution, but also of mutual *complementarity*. The threat of refusing to hedge with a risky dealer is more credible when contagion leads to lose a high final payoff. When a dealer posts collateral in advance, his final payoff in case of success is higher than under a first-best contract. In turn, it is more costly to suffer

[6] A possible extension of the model should consider the possibility to have multiple hedging relationships. In this case, my conjecture is that optimal default fund contributions should be divided among hedging partners in a way proportional to their relative weight in the defaulting member portfolio.

[7] Miglietta et al. (2015) quantifies the cost imposed by initial margins in repo markets.

contagion, and a safe dealer is less willing to trade with a risky counter-party. This behaviour favours other dealers, because it allows an effective implementation of peer monitoring, reducing initial collateral requirements. In other words, peer monitoring is a substitute for collateral, but, at the same time, a dealer benefits from *other* dealers' collateral, because they are more likely to cross-monitor, relaxing ex ante his *own* collateral requirements.

Despite the advantages in terms of peer monitoring, the optimal default waterfall has some drawbacks in terms of systemic risk. In particular, promoting the incentives to isolate a financial dealer may accelerate its insolvency. However, these concerns could be better managed within an effective resolution mechanism, or through emergency lending, rather than favouring its possibility to trade through a CCP. Moreover, exploiting early warning signals from other dealers may allow to cope with a risky dealer at an early stage, mitigating the loss from a later intervention. This concern is particularly important when the financial institution is vastly interconnected with the rest of the financial system.[8]

In the next section I discuss the related literature. Section 3 presents the model setup. Section 4 characterizes the equilibria of the baseline model. Section 5 extends the model to multiple dealers. Section 6 further discusses the results. Section 7 concludes. All proofs are presented in Appendix.

2.2 Related literature

My paper is closest to a recent literature that studies how the insurers' agency problem change with the introduction of a CCP. Several papers stress the possibility to pool and diversify idiosyncratic default risk through a CCP. In this respect, a CCP provides a missing insurance market, and improves the allocation of risk among traders. In Carapella and Mills (2012), clearing services facilitate the possibility to undertake socially valuable transactions; both counter-parties have less incentives to acquire information on the asset value, and, in turn, this common incentive to remain ignorant reduces the negative effects of adverse selection. Antinolfi et al. (2014) also point out that a CCP reduces traders' incentive to acquire information on counter-party risk.[9] However, this lack of screening incentives results in higher collateral requirements, as an instrument to mitigate limited commitment. When the insurance motive is not too strong, buyers may acquire information on counter-party risk, and use clearing services only with the

[8]Incidentally, CCPs may play an important role also for the ordered liquidation of the open positions held by a defaulting member. In particular, fire sales episodes could be greatly reduced thanks to better coordination, and the possibility to auction off the portfolio to other surviving members. However, there is still a lot of scope for improving resolution procedures and the post default auction process. The debate is still underway in the financial community; see *"CCPs confront the difficult maths of default management"*, Risk 28/01/15.

[9]A similar point is informally discussed in Koeppl and Monnet (2012).

most risky traders. Biais et al. (2012b) consider a model in which insurance buyers may exert effort to screen less risky counter-parties, and they compare bilateral versus centralized clearing in the presence of idiosyncratic or aggregate risk. They conclude that central clearing is strictly welfare improving as it allows to diversify idiosyncratic risk. Without aggregate shocks, the allocation provides full insurance, no collateral is posted, and insurance buyers avoids screening costs because, de facto, they cross-insure through the CCP. In the presence of aggregate shocks, insurance sellers are again valuable, and screening may be necessary. For this purpose, the optimal clearing contract only offers partial insurance to maintain buyers' incentives to screen only good counter-parties. Lastly, Koeppl (2013) proposes a new disciplining device, which takes the form of an endogenous 'liquidity' cost. If it is easy to find a new counter-party, insurance buyers can credible threaten not to contract with a risky counter-party. A CCP can lead to a negative feedback effect: an increase in collateral requirements, then a reduction in the amount of profitable transactions, and, lastly, a decrease in market 'liquidity'. In my model, I abstract from the advantages of risk pooling—all players are risk-neutral—and I consider exclusively the incentive role of collateral. In particular, I focus on how to exploit the interactions in the inter-dealer market to economize on collateral. In this respect, a CCP should be designed to maximize the effectiveness of these mechanisms. My modelling of a CCP stresses the incentive role of ex post loss mutualization mechanisms—such as additional default fund contributions or variation margin haircutting—rather than ex ante collateral posting.

My results are also related to several papers on peer monitoring. Rochet and Tirole (1996) argue that the existence of a decentralized inter-bank market can be only motivated by peer monitoring. In turn, to be effective, the incentive to monitor other banks requires to be responsible for losses on inter-bank loans, and not to be protected by a lending of last resource policy. Similarly, in my model peer monitoring requires surviving dealers to be exposed to contagion risk, and to pay additional penalties when they previously hedged with a defaulting member. Stiglitz (1990) discusses the role of penalties as an incentive device for peer monitoring. Focusing on microfinance, Varian (1990) and Ghatak (2000) show how joint liability contracts lead to an endogenous peer selection among borrowers of similar credit risk. In my model, there is an analogous assortative matching outcome. Differently, I study how to design the default waterfall to exploit dealers' information, and I consider to which extent the effectiveness of endogenous matching depends on the correlation among dealers' credit shocks.

Two influential papers discuss other features of CCPs. Duffie and Zhu (2011) point out that CCPs increase netting benefits only if multilateral netting for a single asset class dominates the bilateral netting achievable across different underlying assets. In this respect, they stress the importance not to fragment trade across too many CCPs.

Acharya and Bisin (2014) argue that the lack of transparency in OTC market imposes a 'counter-party risk' externality, leading to excessive leverage and increased default risk. A CCP improves market outcomes because—by concentrating all trade through novation—it provides complete information of trade positions. I abstract from netting and trade transparency, and I focus exclusively on the interaction between the design of the default waterfall and the risk-shifting incentives.

2.3 Baseline model

There are four periods $t = 0,1,2,3$ and two financial dealers F_A, F_B who face a population of competitive, deep pockets, and risk-neutral investors. I use F_i and F_{-i}, $i = A,B$, to denote dealer i and his counter-party, respectively. Each dealer has an identical project which requires an initial outlay of $I > 0$ dollars and pays either a return $R > 0$ (deterministic) or zero. The probability that F_i's investment returns zero at $t = 3$ is a random variable d_i which realizes at $t = 2$. The probability distribution of d_i depends on the effort choice $e_i \in \{0,1\}$ which F_i chooses at $t = 1$. If $e_i = 1$ he incurs a cost $c > 0$ and he receives R at $t = 3$ for sure; if F_i shirks ($e_i = 0$), then d_i is randomly drawn from a distribution on $[0,1]$, with cumulative density function $G(\cdot)$ differentiable a.e., and expected value m. The investment is profitable only if F_i exerts effort, i.e $c < Rm$ and $(1-m)R < I$.

At $t = 2$ the tuple $\mathbf{d} = (d_A, d_B)$ is realized. It is not observable to investors or a judicial court but only to dealers. After observing $\mathbf{d}$ each F_i decides whether he is willing to 'hedge' or not with F_{-i}.[10] If *both* agree they avoid a cost $L > 0$, otherwise each dealer incurs this cost. For example, L can be interpreted as the extra effort required to manage an unhedged position. Hedging avoids the cost L, but it increases the risk of contagion: if at $t = 3$ dealer F_{-i} defaults, F_i may default as well with probability $\gamma \in (0,1]$. The γ parameter captures how risky is the inter-dealer market in transmitting default shocks to other counter-parties. The hedging decision at $t = 2$ is a reduced form to capture any financial relationship aimed at reducing operating costs, but determining a financial exposure towards F_{-i} credit risk. If both dealers agree to hedge and $\mathbf{d} = (d_i, d_{-i})$, F_i defaults at $t = 3$ with probability $d_i + (1-d_i)d_{-i}\gamma$.

At $t = 0$ each dealer F_i simultaneously offers a contract $w_i = (p_i, k_i)$ to outside investors. Contracts $\mathbf{w} := (w_A, w_B)$ are publicly observed. Each F_i receives I from investors, and promises a transfer p_i, to be paid at $t = 3$, and a quantity $k_i \geq 0$ of collateral posted at $t = 0$. Dealers can produce the collateral asset at $t = 0$ at a per unit

[10] I refer to this action between two dealers as 'hedging', but the modelling is extremely stylized. I use this terminology only because derivative dealers mostly use inter-dealer markets for this purpose. For the model results, what is necessary is that inter-dealer markets reduce the costs of running the project initially financed by investors.

cost $\mu > 1$, and at $t = 3$ one unit of collateral is worth one dollar (normalized). The discount rate of all market participants is normalized to one. If F_i defaults, collateral k_i is transferred to investors, while if he honours his obligations F_i gets back the asset. After observing the contract offers $\mathbf{w}$, investors decide whether to accept or reject the offer. I consider $t = 0$ to be divided in two sub-periods: in the first one dealers simultaneously offer contracts, and in the second one investors accept or reject.

Clearly, without a moral hazard problem no collateral asset would be produced since investors are risk neutral. As implicit in the restriction of the contractual space (p_i, k_i), dealers cannot offer contracts contingent on the hedging outcome at $t = 2$. However, the hedging outcome can be ex-post verified—for example during a bankruptcy procedure—after a dealer defaults.[11] The timing of the game is illustrated in Figure 2.1.

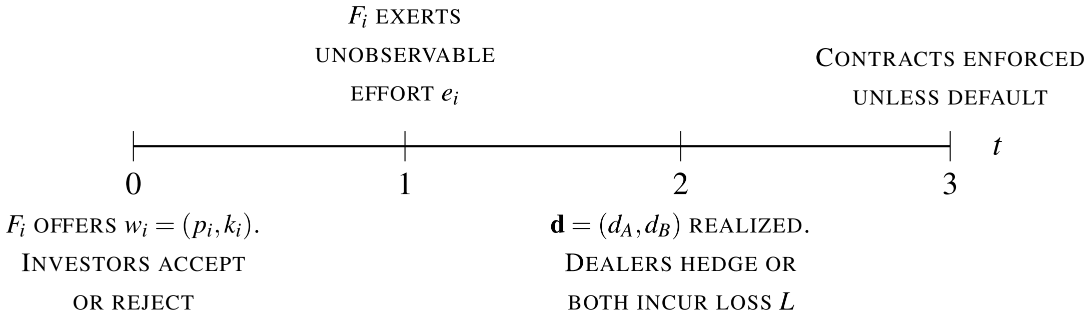

Figure 2.1: Timing of the game.

I consider the subgame perfect equilibria of the game. A strategy profile is subgame perfect if its restriction is a Nash equilibrium in every proper subgame. The proper subgames coincide with each period t and sub-period (for $t = 0$): at $t = 2$ both dealers know $(\mathbf{w}, \mathbf{d})$ and decide whether to hedge or not; at $t = 1$ dealers decide the effort decision, knowing contracts $\mathbf{w}$ and investors' acceptance decision; in the second sub-period of $t = 0$, investors accept contract w_i after observing $\mathbf{w} = (w_A, w_B)$; in the first sub-period of $t = 0$, each F_i simultaneously offers a contract w_i.

2.4 Equilibrium with two dealers

To highlight the effects of peer monitoring, I first analyze the equilibrium contract for a single dealer, i.e. at $t = 2$ the game has no hedging, and dealers do not incur the loss

[11] The possibility to write investors' contracts contingent on the hedging outcome could enlarge the possibilities to implement a first-best contract, because it mitigates the risk-shifting problem. Financial contracts contingent on other contracts are often difficult to implement in practice as they may require real-time information on dealers' overall portfolio position. However, the Lehman's bankruptcy case in 2008 has shown how difficult, lenghty, and open to court litigation it is to evaluate a dealer's portfolio position. Although a trade repository or a CCP may greatly improve these possibilities, a real time monitoring seems a too ambitious goal for the time being.

L. In section 2.4.2 I present the equilibrium with two dealers. Lastly, Section 2.4.3 discusses how the optimal design of penalties changes the equilibrium outcomes.

2.4.1 Autarchic equilibrium

In a first-best allocation F_i invests in the project and exerts effort. Collateral k_i is not used because effort is contractible and no incentive compatibility constraint arises. As one unit of collateral costs $\mu > 1$, it is optimal to only pay in the last period, i.e. $p_i = I$ and $k_i = 0$. The maximum investment level which can be profitably financed is $I^* := R - c$.

For simplicity, I denote the final payoff at $t = 3$ for F_i with $\pi_i := R - p_i + k_i$. If effort is not observable and verifiable, a dealer exerts effort at $t = 1$ only if the contract is incentive compatible

$$\pi_i - c \geq (1-m)\pi_i \qquad \rightarrow \qquad \pi_i \geq \frac{c}{m} \tag{2.1}$$

The first-best contract—$p_i = I$ and $k_i = 0$—is incentive compatible only if

$$I \leq R - \frac{c}{m} = I^* - \frac{1-m}{m}c := I_a^{*}$$ [12]

If $I > I_a^*$ the first-best contract is not implementable and collateral k_i must be used to satisfy equation (2.1). An incentive compatible contract requires to post collateral:

$$k_i \geq \frac{c}{m} - (R - p_i) \tag{2.2}$$

If this is the case, investors accept to finance the project only if $p_i \geq I$. In equilibrium, investors break-even and equation (2.2) is binding. Investing is profitable if and only if:

$$R - I - (\mu - 1)\left(\frac{c}{m} - (R - I)\right) \geq 0 \quad \rightarrow \quad I \leq I^* - \frac{\mu - 1}{\mu}\frac{1-m}{m}c := I_a^k \tag{2.3}$$

Compared to the first-best contract, the equilibrium outcome is suboptimal when: (i) $I \in (I_a^*, I_a^k]$ because costly collateral must be used; or (ii) $I \in (I_a^k, I^*]$ because it is not possible to finance the investment.

[12] The subscript a stands for 'autarchy'.

2.4.2 Peer monitoring equilibrium

I now extend the autarchic model by considering the effect of the 'hedging game' at $t = 2$. In the sub-game starting at $t = 2$, each dealer F_i knows contracts $\mathbf{w}$ and default probabilities $\mathbf{d}$, and he decides whether he is willing to hedge or not with F_{-i}. If F_{-i} is willing to hedge, it is weakly optimal for F_i to accept if and only if:

$$(1-d_i)\pi_i - (1-d_i)d_{-i}\pi_i\gamma \geq (1-d_i)\pi_i - L \quad \rightarrow \quad d_{-i} \leq \frac{L}{(1-d_i)\gamma\pi_i}$$

In words, F_i wants to hedge only if the default probability of the other dealer is sufficiently low. The threshold value negatively depends on the expected loss from contagion—$(1-d_i)\gamma\pi_i$—and positively on L. In particular, F_i is less willing to be financially exposed to F_{-i} when the loss L from refusing to hedge is small, contagion probability γ is significant, its own credit risk d_i is low, and its future payoff π_i is high.

Dealers hedge for sure when they both exert effort, and no loss L is borne. Indeed, F_{-i} is always willing to hedge if F_i is safe ($d_i = 0$) because there is no contagion risk. However, if F_i did not exert effort but F_{-i} did, the probability of hedging is only $G\left(\frac{L}{\gamma\pi_{-i}}\right)$ as F_{-i} may refuse to trade when d_i is too large.

In a first-best contract both dealers exert effort, hedge with probability one, and the equilibrium contract is $p_i = I$ and $k_i = 0$. If effort were contractible the hedging decision would not be useful to provide incentives.

Since effort is not contractible, its provision requires to have an incentive compatible contract. The important novelty in the two dealers model—relative to the autarchic one—is the possibility to use the inter-dealer transaction as an incentive device. If a dealer is too risky, the other dealer may reject to hedge, as the loss from a potential contagion is too high. This potential exclusion from the inter-dealer market at $t = 2$—which I name *threat of ostracism*—may provide additional incentives ex ante, at $t = 1$, to exercise costly effort and avoid the possibility to incur the loss L.

I proceed to characterize the equilibrium. Consider a path in which F_{-i} exercises effort and has a final payoff π_{-i} at $t = 3$. Hence, F_i is always willing to hedge with F_{-i} as the latter is safe. If both dealers undertake the investment, the optimal contract solves:

$$\max_{p_i,k_i} \quad R - p_i - (\mu-1)k_i - c$$

$$s.t. \quad \pi_i = R - p_i + k_i \geq \frac{c - \left(1 - G\left(\frac{L}{\gamma\pi_{-i}}\right)\right)L}{m} \qquad \text{(IC)}$$

$$p_i \geq I \qquad \text{(IR)}$$

Without setting up a Lagrangian, it is easy to realize that the optimal contract always features $p_i = I$. The first-best contract with $k_i = 0$ is still implementable when the IC constraint is slack, i.e.:

$$\begin{aligned} I \leq & \quad R - \frac{c - \left(1 - G\left(\frac{L}{\gamma\pi_{-i}}\right)\right)}{m} L = I^* - \frac{1-m}{m} c + \frac{\left[1 - G\left(\frac{L}{\gamma\pi_{-i}}\right)\right]}{m} L \\ = & \quad I_a^* + \frac{\left(1 - G\left(\frac{L}{\gamma\pi_{-i}}\right)\right)}{m} L := I_c(\pi_{-i}) \end{aligned}$$

If $L < \gamma\pi_{-i}$ the threat not to hedge relaxes the incentive compatibility constraint compared to an autarchic situation as $\left[1 - G\left(\frac{L}{\gamma\pi_{-i}}\right)\right] L > 0$. This is the case when L is not too large, otherwise F_{-i} would always accept to avoid the loss L, even if this could lead to very high contagion risk. In other words, the impact of a greater L is non-monotone. On the one hand, a higher L provides more incentives to exercise effort because it is more costly not to find a counter-party willing to hedge. On the other hand, a higher L reduces the effectiveness of peer monitoring, as a safe dealer may still prefer to avoid the loss L at the cost of a higher contagion risk. This countervailing force is stronger when contagion risk γ is low, as the expression $G\left(\frac{L}{\gamma\pi_{-i}}\right)$ points out. Importantly, the peer monitoring incentives are stronger when the other dealer has a larger stake π_{-i} at $t = 3$.

Tenuous peer monitoring incentives worsen the agency problem, and collateral may be necessary to increase the skin in the game. This is the case for $I > I_c(\pi_{-i})$. Being costly, the amount k_i of collateral posted at $t = 0$ is set just enough to satisfy the IC constraint, i.e.:

$$k_i = \frac{c - \left[1 - G\left(\frac{L}{\gamma\pi_{-i}}\right)\right] L}{m} - (R - I) = I - I_c(\pi_{-i})$$

Substituting this quantity into F_i's profit function, I get the maximum investment cost which can be financed through an incentive compatible contract:

$$R - I - (\mu - 1)\left[\frac{c - \left[1 - G\left(\frac{L}{\gamma\pi_{-i}}\right)\right] L}{m} - (R - I)\right] \geq 0$$

i.e.

$$\begin{aligned} I \leq & \quad I^* - \frac{\mu - 1}{\mu m}\left[(1-m)c + \left[1 - G\left(\frac{L}{\gamma\pi_{-i}}\right)\right] L\right] \\ = & \quad I_c(\pi_{-i}) + \frac{1}{\mu m}\left[(1-m)c - \left[1 - G\left(\frac{L}{\gamma\pi_{-i}}\right)\right] L\right] := I_c^k(\pi_{-i}) \end{aligned} \tag{2.4}$$

Proposition 2.4.1 characterizes the unique subgame perfect equilibrium of the game.

For this purpose, it is useful to introduce two quantities:

$$\begin{aligned} \xi_c(\pi) &:= \frac{c - \left[1 - G\left(\frac{L}{\gamma\pi}\right)\right] L}{m} \\ \mathscr{C}(\pi) &:= \frac{c}{\left[1 - G\left(\frac{L}{\gamma\pi}\right)\right] L} \end{aligned} \tag{2.5}$$

The first quantity $\xi_c(\pi)$ is the minimum final payoff that dealer F_i has to receive in order to satisfy his incentive compatibility constraint, given that the other dealer F_{-i} has final payoff π. The second expression $\mathscr{C}(\pi)$ is the ratio between effort cost, and the expected loss incurred when a shirking dealer does not find a counter-party willing to hedge.

Proposition 2.4.1 *The game has a unique perfect subgame equilibrium. Dealers offer identical contracts: $p_i = p$ and $k_i = k$ for $i = A, B$.*

- *The first-best contract $p = I, k = 0$ is implementable for $I \leq \min\{I^*, I_c^*\}$ where I_c^* is the unique solution to*

$$R - I = \xi_c(R - I) \tag{2.6}$$

  Moreover, $I_c^* \geq I^*$ if and only if $\mathscr{C}(R - I) \leq \frac{1}{m}$.

- *If $\mathscr{C}(R - I) > \frac{1}{m}$ and the investment is undertaken, dealers final payoff π_c^* is the solution to $\pi = \xi_c(\pi)$. The second-best contract $p = I, k = \pi_c^* - (R - I) > 0$ is profitable only if $I \in \left(I_c^*, I_c^k\right]$, where*

$$I_c^k = I^* - \frac{\mu - 1}{\mu m}\left[(1 - m)c - \left[1 - G\left(\frac{L}{\gamma\pi_c^*}\right)\right] L\right] = I^* - \frac{\mu - 1}{\mu}(\pi_c^* - c) \tag{2.7}$$

- *No investment is undertaken for $\left(I_c^k, I^*\right]$.*

Compared to an autarchic equilibrium, the hedging game introduces the possibility to exploit dealers' superior information on credit risk. Hedging reduces operating costs, but it also creates an incentive not to trade with risky counter-parties because of contagion. Ex ante, this acts as a commitment device to exert effort, and it reduces the need to use collateral as a way to implement proper incentives. Peer monitoring is helpful to implement the first-best contract for a larger set of investment costs only if $\left[1 - G\left(\frac{L}{\gamma\pi_c^*}\right)\right] L > 0$, i.e $L < \gamma\pi_c^*$. Since $\pi_c^* \geq R - I$, a sufficient condition is $L < \gamma(R - I)$. Even when collateral is used to restore incentives, for $I \in \left(I_c^*, I_c^k\right]$, peer monitoring incentives further alleviate the risk-shifting problem. Substituting π_c^* in equation (2.7) and comparing it with equation (2.3), it is straightforward to notice that $I_c^k > I_a^k$ since $G\left(\frac{L}{\gamma\pi_c^*}\right) < G\left(\frac{L}{\gamma(R-I)}\right) < 1$ as $\pi_c^* > R - I$. The last observation points out an additional role played by collateral: by increasing the final payoff at $t = 3$, it enhances peer monitoring incentives. Indeed, a dealer is more likely to reject a hedge when the potential counter-party is risky, as he risks a larger loss in case of contagion. In comparison to a first-best contract, a safe dealer is more likely to reject a hedge from a risky dealer.

Nonetheless, when the cost L of refusing to hedge is too high, the incentives to peer monitor are less credible. A safe dealer has less incentive to ostracize a risky one if no other hedge is possible; the immediate loss L may be larger than the expected loss from contagion risk. In this situation, peer monitoring does not alleviate the agency problem and the optimal contract is identical to the one under autarchy. Similarly, if the financial assets exchanged in the inter-dealer market have low contagion risk (low γ), dealers find optimal to hedge also with a risky counter-party as the risk to incur a loss is low. For example, collateralized loans such as repos may reduce contagion risk, but they may inhibit peer monitoring incentives. Rochet and Tirole (1996) highlight the importance to have sufficient losses on the inter-bank market in order to create sufficient peer monitoring incentives. In my setup, when peer monitoring incentives are low—i.e. $\mathscr{C}(\pi)$ high—there is an increase in collateral requirements, and a higher cost of capital. In section 2.4.3, I explain how a CCP loss mutualization mechanism may help to solve this difficult trade off between peer monitoring incentives and contagion risk.

2.4.3 Optimal CCP loss mutualization design

In this section I introduce a market infrastructure, such as a CCP, with the possibility to: (i) impose an initial margin requirement k in the form of collateral; (ii) upon the default of a dealer, require additional payments from the other dealer (if still solvent) in order to cover losses suffered by outside investors. In a broader sense, I interpret a CCP as an overarching contract at $t = 0$ between F_A and F_B which introduces a system of *penalties* for the default of the other dealer. These transfers have no insurance role—investors are risk-neutral—but they may be used to incentivize the threat of ostracism. As I am going to explain, this mechanism is successful only if these penalties depend on the hedging decision at $t = 2$.

In real world practice, a CCP default waterfall includes the possibility to call for additional contributions to the default fund. These transfers from dealers to the CCP are usually limited in value, and they do not depend on previous trading patterns among CCP members. I consider a slightly more general situation, in which additional default funds may depend on whether dealers hedge at $t=2$. If F_{-i} defaults, dealer F_i has to pay non-negative transfers τ_1 or τ_0 to the CCP, depending on whether they hedged or not, respectively.

I analyze the impact of introducing τ_0 and τ_1 on the equilibrium outcomes, and I assume the CCP sets contracts $\mathbf{w}$ and penalties $\tau_j,\ j=0,1$. Leaving the contract decision w_i up to dealers, and letting the CCP only set τ_j, would only complicate the analysis without changing final outcomes. I derive the optimal (τ_0,τ_1) in two steps. First, I characterize the equilibrium for every feasible (τ_0,τ_1) and final payoff π_{-i} for the other dealer. Second, I show that it is optimal to increase τ_1 as much as possible and to set τ_0 equal to zero.

At $t=2$ a dealer F_i is willing to hedge only if:

$$(1-d_i)\pi_i-(1-d_i)d_{-i}[\gamma h_i+(1-\gamma)\tau_1]\geq(1-d_i)\pi_i-L-(1-d_i)d_{-i}\tau_0$$

Simplifying and rearranging:

$$L\geq(1-d_i)d_{-i}[\gamma\pi_i+(1-\gamma)\tau_1-\tau_0]$$

If $\gamma\pi_i+(1-\gamma)\tau_1-\tau_0<0$ the dealer would always want to hedge. If the opposite inequality holds, the acceptance rule becomes:

$$d_{-i}\leq\frac{L}{(1-d_i)(\gamma\pi_i+(1-\gamma)\tau_1-\tau_0)}=\frac{L}{(1-d_i)h_i}\tag{2.8}$$

As before, if F_{-i} exerts effort but F_i does not, both dealers hedge at $t=2$ with probability $G\left(\frac{L}{h_{-i}}\right)$. Assuming F_i expects F_{-i} to pay the effort cost, the optimal contract solves:

$$\begin{aligned}
\max_{p_i,k_i}\quad & R-p_i-(\mu-1)k_i-c \\
s.t.\quad & R-p_i+k_i\geq\frac{c-\left[1-G\left(\frac{L}{h_{-i}}\right)\right]L}{m}:=\xi_\tau(\pi_{-i},\tau_0,\tau_1) && \text{(IC)}\\
& p_i\geq I && \text{(IR)}
\end{aligned}$$

If $h_{-i}\leq 0$, F_i always accepts to hedge, peer monitoring is never effective, and the incentive compatible contract is identical to the autarchic one; see section 2.4.1. If $h_{-i}>0$ peer monitoring mitigate the agency problem when $\frac{L}{h_{-i}}<1$. In this case, the optimization problem is analogous to the one in section 2.4.2, and the second-best contract has $p_i=I$ and:

$$k_i=\max\{0,\ \xi_\tau(\pi_{-i},\tau_0,\tau_1)-(R-I)\}$$

The first-best contract is implementable if and only if $k_i \leq 0$, i.e.

$$I \leq R - \xi_\tau(\pi_{-i}, \tau_0, \tau_1) = R - c - \frac{1}{m}\left[(1-m)c - \left[1 - G\left(\frac{L}{h_{-i}}\right)\right]L\right] := I^*(\pi_{-i}, \tau_0, \tau_1)$$

If the first-best is not attainable, the second-best contract must use collateral to satisfy the incentive compatibility constraint. The contract is profitable only if the investment I satisfies:

$$I \leq R - \xi_\tau(\pi_{-i}, \tau_0, \tau_1) = R - c - \frac{\mu - 1}{\mu m}\left[(1-m)c - \left[1 - G\left(\frac{L}{h_{-i}}\right)\right]L\right] := I^k(\pi_{-i}, \tau_0, \tau_1) \quad (2.9)$$

The equilibrium is symmetric and it differs from the one in Proposition 2.4.1 by the presence of τ_0 and τ_1 in h_{-i}. In turn, this quantity changes $G(\cdot)$ and it can positively or negatively affect peer monitoring incentives.

If the incentive compatibility constraint is binding, the equilibrium payoff $\pi_i = \pi(\tau_0, \tau_1)$, $i = A, B$, is the unique solution to the equation:

$$\pi = \xi_\tau(\pi, \tau_0, \tau_1)$$

The maximum investment level $I^*(\tau_0, \tau_1)$ implementable with a first-best contract solves the equation $R - I = \xi_\tau(R - I, \tau_0, \tau_1)$. By implicit differentiation, it is:

$$\frac{\partial I^*(\tau_0, \tau_1)}{\partial \tau_0} < 0 \qquad\qquad \frac{\partial I^*(\tau_0, \tau_1)}{\partial \tau_1} > 0$$

if the derivative is evaluated at a point such that $G\left(\frac{L}{\gamma(R-I)+(1-\gamma)\tau_1-\tau_0}\right) < 1$. It is optimal to increase τ_1 as much as possible and to set $\tau_0 = 0$. In order to be feasible, the penalties imposed on F_i cannot be larger than the final payoff at $t = 3$, i.e. $\tau_j \leq \pi$, $j = 0, 1$. As a result, peer monitoring incentive are exploited as much as possible by the penalties $\tau_0 = 0$ and $\tau_1 = R - I$.

The highest investment cost implementable with collateral—$I^k(\tau_0, \tau_1)$—is increasing in $h = \gamma\pi + (1-\gamma)\tau_1 - \tau_0$; see equation (2.9). An increase in τ_1 (τ_0) directly increases (decreases) h but, indirectly, it has the opposite effect through π, which solves $\pi = \xi_\tau(\pi, \tau_0, \tau_1)$. Computing the total derivative with respect to h, it is immediate to show that the direct effect dominates on the indirect one:

$$\begin{aligned}
\frac{\partial h}{\partial \tau_0} &= \gamma\frac{\partial \pi}{\partial \tau_0} - 1 = \frac{\gamma}{\gamma + \frac{mh^2}{Lg(h)}} - 1 < 0 \\
\frac{\partial h}{\partial \tau_1} &= \gamma\frac{\partial \pi}{\partial \tau_1} + (1-\gamma) = (1-\gamma)\left(1 - \frac{\gamma}{\gamma + \frac{mh^2}{Lg(h)}}\right) > 0
\end{aligned}$$

Therefore, the optimal CCP design leads to a payoff π^*_τ which solves $\pi = \xi_\tau(\pi, 0, \pi)$, and to a maximum investment level $I^k(0,\pi^*_\tau)$. It can be alternatively expressed as:

$$I^k(0,\pi^*_\tau) = R - c - \frac{\mu-1}{\mu}[\pi^*_\tau - c]$$

As $\xi_c(\pi) > \xi_\tau(\pi,0,\pi)$ for every $\pi > 0$, it follows that $\pi^*_\tau < \pi^*_c$. By equation (2.7), it is immediate to conclude that $I^k(0,\pi^*_\tau) > I^k_c$. In words, the optimal CCP design increases the 'pledgeable income' compared to a bilateral peer monitoring equilibrium.

The previous results suggest how a well designed loss mutualization scheme may promote peer monitoring and alleviate the moral hazard problem. The CCP has to punish as much as possible ($\tau_1 = \pi^*$) a dealer who traded with a later defaulting party, while it should not punish a dealer who refused to trade ($\tau_0 = 0$). This system of penalties: (i) reduces contagion risk at $t = 2$ as it discourages to trade with a risky counter-party; and (ii) at $t = 1$ it provides additional incentives to exert effort and be able to hedge at $t = 2$. In this model with only two dealers, the inter-dealer market provides incentives to avoid being risky by a threat of ostracism from a safe counter-party. After a deviation from the equilibrium path, a dealer faces a safe counter-party and the CCP structure enhances his incentives to refuse a hedge. In a market without CCP, the cost of hedging with a risky dealer is the possibility to lose future profits in case of contagion; a CCP imposes additional penalties as it expropriates profits ($\tau_1 = \pi^*_\tau$) even when there is no contagion.

The penalties enhance the effectiveness of peer monitoring. In a market without CCP peer monitoring is effective only if $L < \gamma(R-I)$, while the optimal CCP design ($\tau_0 = 0, \tau_1 = \pi^*_\tau$) makes it relevant if $L < R - I$. It is no longer important to have a higher contagion probability γ, because τ_1 achieves the same incentive effects—expropriating F_i's payoff at $t = 3$—without relying on the extreme punishment of a default due to contagion. This is a clear advantage because F_i's contagion also harms his investors. As a consequence, an optimal market architecture should impose a well designed CCP and reduce γ as much as possible. However, the model abstracts from the determinants of γ and the possible costs related to a decrease in contagion risk throurgh higher collateralization. In this case, a standard cost-benefit analysis should apply to determine the optimal trade-off.

The previous results point out the importance to condition penalties on a measure of previous trade with the defaulting dealer. If penalties are always imposed on a surviving member—independently from the hedging decision—then $\tau_0 = \tau_1 = \tau$. In this case $\pi(\tau,\tau)$ solves:

$$\pi = \frac{c - \left[1 - G\left(\frac{L}{\gamma(\pi-\tau)}\right)\right]L}{m}$$

The maximum investment levels are $I^*(\tau,\tau) < I_c^*$ and $I^k(\tau,\tau) < I_c^k$ for every $\tau > 0$, and the negative impact on the investment opportunities is larger the higher is τ. The intuition is very simple: if F_i has to always cover F_{-i} losses (even partially), it is more convenient to trade and avoid the extra hedging cost L. Moreover, the risk of contagion is less damaging as it avoids to pay for F_{-i} losses.

2.5 Equilibrium with multiple dealers

I extend the baseline model to include multiple dealers. The main goal is to understand how the number of participants in the inter-dealer market may change the incentive to exercise effort, and whether the optimal CCP design should change relative to the model with two dealers. In real world markets, dealers hold multiple positions with other dealers. However, I abstract from the issue of credit risk diversification among multiple dealers, and I continue to assume dealers hedge in bilateral relationships.[13] This assumption streamlines the exposition of the economic forces at work when dealers *sort* based on credit risk.

In the baseline model, dealers do not choose their hedging counter-party because the inter-dealer market has only two members. After observing the default probability, they only decide whether to accept a hedge or not. Moreover, in the model in section 2.4, exerting effort leads to a zero default probability and, under the optimal contract, all dealers survive at $t = 3$. As a result, if a dealer shirks at $t = 1$, he expects to face a safe dealer at $t = 2$. Extending this framework to multiple dealers would not change the interactions of the hedging game, and the final outcome would be identical to the one in section 2.4.2. In other words, the number of dealers is irrelevant if exerting effort leads all dealers to be always safe. In this case, peer monitoring is effective through the threat of being *ostracized* when too risky.

If exerting effort does not guarantee to be safe, but it may lead, occasionally, to be risky, a new economic mechanism arises. It provides additional incentives to avoid risk-shifting, especially when the number of dealers in the market is sufficiently large. In this extension, exerting effort continues to reduce the threat of ostracism, but it also increases the chances to hedge with a safe counter-party. In the next subsections, I am going to explain this endogenous peer selection mechanism, and how it changes the incentives to exercise effort and, in turn, final outcomes.

[13] Although I do not model multiple hedging relationships, it seems reasonable to conjecture that losses should be distributed among counter-parties proportionally to their relative share of trade with the defaulting dealer. Modelling multiple trade relationships among dealers, and developing appropriate measures of dealers' interconnectedness—in relation to the peer monitoring problem—are interesting future research venues.

2.5.1 Model

I extend the baseline model in the simplest possible manner to preserve analytical tractability. Assume there are N (even number) identical dealers at $t=0$. At $t=2$ dealers can either be *safe* or *risky*. If safe, a dealer survives for sure at $t=3$ unless he defaults by contagion. If risky, at $t=3$ he defaults with probability d, and survives with probability $1-d$ (unless contagion). The random variable d comes from a differentiable cumulative distribution $G(\cdot)$, and it is identical for all risky dealers at $t=2$. At $t=1$, exerting effort leads with probability $1-\alpha$ to be safe, and with probability α to be risky; if a dealer shirks, he is risky with probability one. Conditional on N dealers exerting effort, the probability to have $l \leq N$ safe dealers comes from a correlated binomial distribution with probability distribution function $\mathbb{P}_N(l)$. I assume the random variable d to be independent from the realization l of safe dealers.

As in the baseline model, dealers are the only ones to observe the default probability of other dealers. Once they observe that l dealers are safe and $N-l$ are risky, they simultaneously choose to match with another dealer. I do not model explicitly this matching process, and I directly consider the following outcome:

- If l is even, every dealer matches with a dealer of identical credit risk: safe with safe, and risky with risky.
- If l is odd, each safe (risky) dealer matches with probability $\frac{l-1}{l}$ $(\frac{N-l-1}{N-l})$ to a safe (risky) dealer, and with probability $\frac{1}{l}$ $(\frac{1}{N-l})$ with a risky (safe) dealer.[14]

I consider this matching outcome reasonable, on the grounds that there is a mutual benefit for safe types to hedge with each other, and avoid contagion risk. The final outcome leads to *assortative matching*, unless there are not enough pairs of the same type and, for each group, one dealer at random has to trade with a dealer of different credit risk. From an individual point of view, this probability of 'mismatch' is lower the higher is the number N of dealers in the market.

Let $j \in \{gg, gb, bg, bb\}$ denote the possible matching pairs for a dealer, where the first letter indicates his credit risk (g for safe, b for risky), and the second letter the one of his trading partner. If all N dealers exert effort, the probability p_j at $t=1$ is:

$$
\begin{aligned}
p_{gg} = &\ \sum_{l=0}^{N} \mathbb{P}_N(l)\frac{l}{N}\left[\mathbb{I}_{\{1\text{ even}\}} + \left(1-\frac{1}{l}\right)\mathbb{I}_{\{1\text{ odd}\}}\right] = 1-\alpha-\frac{1}{N}\sum_{l=0}^{N}\mathbb{P}_N(l)\mathbb{I}_{\{1\text{ odd}\}} \\
p_{bb} = &\ \sum_{l=0}^{N} \mathbb{P}_N(l)\frac{N-l}{N}\left[\mathbb{I}_{\{1\text{ even}\}} + \frac{N-l-1}{N-l}\mathbb{I}_{\{1\text{ odd}\}}\right] = \alpha-\frac{1}{N}\sum_{l=0}^{N}\mathbb{P}_N(l)\mathbb{I}_{\{1\text{ odd}\}} \\
p_{bg} = &\ p_{gb} = \frac{1}{N}\sum_{l=0}^{N}\mathbb{P}_N(l)\mathbb{I}_{\{1\text{ odd}\}}
\end{aligned}
$$

[14] These are the probabilities, for a single dealer, to match with a safe or risky counter-party. Overall, all dealers' pairs include dealers of the same credit risk when l is even, and there is only one 'mismatched' pair if l is odd.

For the incentive compatibility constraint, another relevant quantity is the probability q_{bg} that, after shirking, a risky dealer matches with a safe dealer, assuming the other $N-1$ dealers exerted effort. It is equal to:

$$q_{bg} = \sum_{l=0}^{N-1} \mathbb{P}_{N-1}(l) \frac{1}{N-l} \mathbb{I}_{\{1 \text{ odd}\}}$$

For example, consider two extreme cases: statistical independence and perfect correlation. In the former, becoming risky is an idiosyncratic shock. In the latter, credit statuses are perfectly correlated, and the shock has an aggregate nature; for example, they may be the result of a deterioration in some macroeconomic variable. For N large, different correlation assumptions do not affect much p_j, while they may continue to matter for q_{bg}. For $N \to +\infty$, the probability q_{bg} is equal to zero in the case of statistical independence, and to $1-\alpha$ if shocks are perfectly correlated. This difference is going to play a relevant role in the incentive compatibility constraint.

2.5.2 Incentive compatibility for N dealers

As in section 2.4.3, a CCP can impose penalties τ_1 and τ_0 on a dealer after the default of his trading partner, and these transfers may depend on the hedging decision at $t=2$. The situation without a CCP is a particular case ($\tau_0 = \tau_1 = 0$). I exclude the possibility to impose penalties on a dealer, based on the default of a dealer other than his hedging partner.[15] In this respect, I model a situation of *bilateral* rather than centralized clearing. I make the assumption that a hedging partner can be identified even if no hedging takes place. This is not an issue when $N=2$, but with more than one pair it is less realistic to impose penalties on dealers who refused to trade. I keep this assumption, although it will be later clear how the optimal CCP design does not rely on it.

To streamline exposition, I restrict attention to equilibria in which dealers of the same credit risk always decide to hedge. This restriction is irrelevant when both dealers are safe, while for risky dealers it requires $d(1-d) \leq \frac{1}{4} \leq \frac{L}{h_i}$, for every $d \in [0,1]$, where $h_i = \gamma\pi_i + (1-\gamma)\tau_1 - \tau_0$. Thanks to this assumption, dealers always hedge, unless a matched pair includes a safe and a risky dealer. In this case, a safe dealer F_{-i} accepts to hedge only if the default probability d is below $\frac{L}{h_{-i}}$ (see equation (2.8)).

For a given $\pi = (\pi_i, \pi_{-i})$ and $\tau = (\tau_0, \tau_1)$—which jointly determine h_i and h_{-i}—the expected payoff from exerting effort at $t=1$, conditional on all other $N-1$ dealers

[15] It is easy to realize that this possibility does not relax the incentive compatibility constraint, but quite the opposite as it would reduce the final expected payoff π_i without any connection to F_i's actions.

exerting effort and having payoff h_{-i}, is:[16]

$$\begin{aligned}\mathbb{E}_{e_i=1}[u_i|\pi,\tau] = &\; p_{gg}\pi_i + p_{bb}(1-m)[(m(1-\gamma)(\pi_i-\tau_1)+(1-m)\pi_i]\\ &+\; p_{bg}\left\{G\left(\tfrac{L}{h_{-i}}\right)\mathbb{E}\left[1-d_i|d_i\le \tfrac{L}{h_{-i}}\right]\pi_i + \left[1-G\left(\tfrac{L}{h_{-i}}\right)\right]\left(\mathbb{E}\left[1-d_i|d_i\ge \tfrac{L}{h_{-i}}\right]\pi_i - L\right)\right\}\\ &+\; p_{gb}\left\{G\left(\tfrac{L}{h_i}\right)\left[\mathbb{E}\left[1-d_{-i}|d_{-i}\le \tfrac{L}{h_i}\right]\pi_i + \mathbb{E}\left[d_{-i}|d_{-i}\le \tfrac{L}{h_i}\right](1-\gamma)(\pi_i-\tau_1)\right]\right.\\ &+\; \left.\left[1-G\left(\tfrac{L}{h_i}\right)\right]\left(\pi_i - L - \mathbb{E}\left[d_{-i}|d_{-i}\ge \tfrac{L}{h_i}\right]\tau_0\right)\right\} - c\end{aligned}$$

The expected payoff at $t=1$ includes all the possible future equilibrium outcomes of the game. With probability p_{gg} dealer F_i is safe, and matches with another safe dealer. They hedge, avoiding the extra cost L, and enjoy for sure their final profit π_i at $t=3$. If instead F_i is safe, and matches with a risky dealer—an outcome whose probability is p_{gb}—he accepts to hedge only if $d\le \frac{L}{h_i}$. If they hedge, he enjoys his final profit π_i if F_{-i} does not default (with expected prob. $\mathbb{E}\left[d_{-i}|d_{-i}\le \frac{L}{h_i}\right]$), or if F_{-i} defaults and F_i does not suffer from contagion (with prob. $\mathbb{E}\left[d_{-i}|d_{-i}\le \frac{L}{h_i}\right](1-\gamma)$). In the last case, F_i has to pay the penalty τ_1 to the CCP. If F_i refuses to hedge, he incurs an immediate loss L but he is not exposed to contagion risk, and enjoys π_i for sure; however, if F_{-i} defaults he might still have to pay a penalty τ_0. If F_i turns out to be risky, he can either match with another risky dealer (with prob. p_{bb}), and avoid the cost L for sure; then, if F_i survives at $t=3$, he enjoys π_i either when F_{-i} does not default (with prob. $1-m$), or when it is not affected by contagion (with prob. $m(1-\gamma)$); nonetheless, in the last case he has to pay τ_1. Finally, if F_i turns to be risky, and matches with a safe dealer (with prob. p_{bg}), he enjoys π_i if he survives at $t=3$, but he incurs an extra cost L at $t=2$ if F_{-i} refuses to hedge (with prob. $1-G\left(\frac{L}{h_{-i}}\right)$).

Simplifying the expression, it is:

$$\begin{aligned}\mathbb{E}_{e_i=1}[u_i|\pi,\tau] = &\; \pi_i(1-m\alpha) - \left[p_{gb}\int_0^{\frac{L}{h_i}} xg(x)\,\mathrm{d}x + m(1-m)p_{bb}\right](\gamma\pi_i + (1-\gamma)\tau_1)\\ &+\; \left[p_{gb}\left(1-G\left(\tfrac{L}{h_i}\right)\right) + p_{bg}\left(1-G\left(\tfrac{L}{h_{-i}}\right)\right)\right]L - \tau_0 p_{gb}\int_{\frac{L}{h_{-i}}}^{1} xg(x)\,\mathrm{d}x - c\end{aligned} \tag{2.10}$$

If F_i decides to shirk ($e_i=0$), he is risky for sure at $t=2$ and his expected payoff at $t=1$ is:

$$\mathbb{E}_{e_i=1}[u_i|\pi,\tau] = q_{bg}\left[(1-m)\pi_i - \left(1-G\left(\frac{L}{h_{-i}}\right)\right)L\right] + (1-q_{bg})(1-m)[(1-m\gamma)\pi_i - m(1-\gamma)\tau_1] \tag{2.11}$$

If F_i matches with a safe dealer, he is not exposed to contagion risk, but he may not hedge if $d>\frac{L}{h_{-i}}$. If both dealers are risky, they hedge for sure but F_i may end up paying a penalty τ_1 if F_{-i} defaults, and he survives.

[16] The expectation operator is computed with respect to the probability measure induced by $G(\cdot)$.

The contract at $t = 0$ is incentive compatible if $\mathbb{E}_{e_i=1}[u_i|\pi,\tau] \geq \mathbb{E}_{e_i=0}[u_i|\pi,\tau]$. Rearranging the expression conveniently, it is equivalent to:

$$\pi_i \geq \frac{c - \kappa_L L - (1-\gamma)\kappa_{\tau_1}\tau_1 + \kappa_{\tau_0}\tau_0}{m(1-\alpha) + \gamma\kappa_{\tau_1}} \tag{2.12}$$

$$\begin{aligned}
\kappa_L &= (q_{bg} - p_{bg})\left[1 - G\left(\frac{L}{h_{-i}}\right)\right] - p_{gb}\left[1 - G\left(\frac{L}{h_i}\right)\right] \\
\kappa_{\tau_1} &= m(1-m)(1 - q_{bg} - p_{bb}) - p_{bg}\int_0^{\frac{L}{h_i}} xg(x)\,\mathrm{d}x \\
\kappa_{\tau_0} &= p_{gb}\int_{\frac{L}{h_{-i}}}^{1} xg(x)\,\mathrm{d}x
\end{aligned}$$

In the remainder of the paper, I am going to focus on the limit case $N \to +\infty$. The asymptotic case better highlights the economic mechanisms at work, and the role of correlation among dealers' shocks.

2.5.3 Equilibrium for $N \to +\infty$

The incentive compatibility constraint in equation (2.12) is complicated by the presence of p_{bg} and p_{gb}, and the terms they multiply. However, the probability of a mismatch is rapidly decreasing in N, and it is not a meaningful economic force at work. Indeed, it arises as an unlikely circumstance because of the restriction to bilateral hedging. As $\sum_{l=0}^{N} \mathbb{P}_N(l) = 1$, it is immediate to realize that:

$$\frac{1}{N}\sum_{l=0}^{N} \mathbb{P}_N(l)\mathbb{I}_{\{1 \text{ odd}\}} \leq \frac{1}{N}$$

As a consequence, for $N \to +\infty$ it holds $p_{gg} \to 1-\alpha$, $p_{bb} \to \alpha$ and $p_{bg} = p_{gb} \to 0$. Intuitively, in equilibrium it is extremely likely to match with a dealer of identical credit risk, irrespective of the correlation among dealers' 'involuntary' credit risk shocks. To understand why correlation does not matter, let's consider two extreme assumptions: (i) statistical independence; (ii) perfect correlation. First, in both cases the probability to get mismatched with a dealer of different credit risk is proportional to $\frac{1}{N}$, hence it is very small for N large. In case (i), at $t = 2$ approximately $(1-\alpha)N$ dealers are safe and αN dealers risky; the small probability of mismatch leads all dealers to hedge with a dealer of equal credit risk. In case (ii), if all dealers exercise effort, they all have identical credit risk at $t = 2$, and they hedge with probability one.

In a first-best, all dealers exert effort, and find almost surely a counter-party of equal credit risk with whom they hedge. As a result, exerting effort leads to a survival probability at $t = 3$ equal to:

$$s := 1 - \alpha + \alpha(1-m)(1-m\gamma) = 1 - \alpha m[1+\gamma(1-m)]$$

i.e. the sum of the probabilities to be safe and match with another safe dealer $(1-\alpha)$, and to be risky (α), hedge almost surely with a risky dealer, and avoid to default either individually $(1-m)$, or through contagion in case of default of the other counter-party $(1-m\gamma)$. Therefore, at $t = 0$ the first-best contract solves:

$$\begin{aligned} \max_{p} \quad & s(R-p) - c \\ \text{s.t.} \quad & sp \geq I \end{aligned} \tag{2.13}$$

The optimal contract has $p = I/s$ and the project is always financed if it has positive net present value (NPV), i.e. only if $I \leq sR - c$. Analogously to section 2.4, I denote with $I^* = sR - c$ the maximum investment level. I assume shirking leads to a negative NPV.

I turn to consider the case in which effort is not observable, nor contractible. For $N \to +\infty$, the incentive compatibility constraint in equation (2.12) becomes:

$$\pi_i \geq \frac{c - q_{bg}\left[1 - G\left(\frac{L}{h_{-i}}\right)\right]L - m(1-m)(1-\alpha-q_{bg})(1-\gamma)\tau_1}{m(1-\alpha)+\gamma m(1-m)(1-\alpha-q_{bg})} := \xi_m(q_{bg}, \pi_{-i}, \tau_0, \tau_1) \tag{2.14}$$

The expression in (2.14) points out two separate economic forces—related to the functioning of the inter-dealer market—that decrease the convenience of shirking, relaxing the incentive compatibility constraint relative to the autarchic equilibrium.

1. *Threat of ostracism.* The quantity $q_{bg}\left[1 - G\left(\frac{L}{h_{-i}}\right)\right]L$ captures the expect loss from shirking due to the possibility of not being able to hedge. It depends on: (i) the probability q_{bg} to match with a safe dealer after shirking; (ii) the probability that a safe dealer rejects to trade as $d > \frac{L}{h_{-i}}$; and (iii) the loss L from not hedging. In particular, the higher is the probability to match with a safe dealer, the more relevant is this economic mechanism. The baseline model of section 2.4 highlights very clearly this threat of ostracism because, in equilibrium, the probability to match with a safe dealer after a deviation is one. As previously discussed, the effect of a larger L is ambiguous: it increases the loss L, but, at the same time, it decreases the chances that a safe dealer refuses to trade, since he is also going to suffer a high loss without a hedge.

2. *Higher costs of a risky hedge.* With infinitely many dealers, matches are not exogenous as in the two dealers' case. In equilibrium, a dealer matches almost surely with another dealer of identical credit quality. Risky dealers are more

likely to default, so their trading partners are more exposed to contagion risk and, possibly, to additional penalties τ_1. Without a CCP, it is convenient not to match with a risky counter-party to avoid contagion risk. This element is captured by the term $\gamma m(1-m)q_{bg}$ in the denominator. A CCP may further increase these incentives to avoid risky counter-parties. To reduce the chances to pay τ_1, it is more convenient to exercise effort, as it is more likely at $t = 2$ to be safe and trade with a safe counter-party. The importance of this channel depends on the quantity $m(1-m)(1-\alpha-q_{bg})(1-\gamma)\tau_1$, which is the expected cost from a 'risky match' after shirking. It depends on: (i) the difference between the probability of matching with a risky dealer when F_i shirks $(1-q_{bg})$ or not (α); (ii) the expected probability that F_i survives, but the hedging partner defaults $m(1-m)(1-\gamma)$; and (iii) the penalty τ_1 that a dealer has to pay in case of default of his counter-party.

The relative importance of these two forces depends—among other parameters—on the probability q_{bg} to match, after a deviation, with a safe dealer. A higher value of q_{bg} increases the weight assigned to the event that a safe dealer refuses to hedge, and it reduces the probability to match with a risky dealer, and possibly pay τ_1 at $t = 3$. The correlation among dealers' shocks influences q_{bg}. Suppose a dealer deviates from an equilibrium in which every dealer exerts effort. At $t = 2$ he is for sure risky, and it is going to match with another dealer. If dealers' 'involuntary' credit risk shocks are independent, then $q_{bg} \to 0$ as N increases; in fact, there are always going to be risky dealers, and the probability of a 'mismatch' is low even after a deviation. In contrast, if there is perfect correlation, it is $q_{bg} = 1-\alpha$, as all the other $N-1$ dealers have the same credit risk at $t = 2$, and their marginal probability to be safe is $1-\alpha$. Intermediate correlation stands between these two extremes.

In section 2.4, the penalties τ_1 or τ_0 do not enter neither the objective function, nor the investors' individual rationality constraint, since in equilibrium all dealers exert effort and no default occurs. However, in this slightly more general setup, default is possible even when dealers exert effort. As a result, τ_1 appears as an additional cost in F_i objective function—to be paid to F_{-i}'s clients conditional on the latter default—and it also shows up as an additional payment that F_i' clients receive from F_{-i}, if the former defaults but the latter does not. Both events have probability $z := \alpha m(1-m)(1-\gamma)$.

In equilibrium, all dealers always hedge irrespective of their credit risk, and the penalty τ_0 does not appear in the profit function or in the budget constraint, although it may still affect the incentive compatibility constraint through h_{-i} (see equation (2.14)). I implicitly assume that τ_1 and τ_0 satisfy the feasibility constraints $\tau_j \leq \pi_i = R - p_i + k_i$,

$j = 1,2$ for all dealers $i \in \{1,...,N\}$. The optimal incentive compatible contract solves:

$$\begin{aligned} \max_{p_i,k_i} \quad & s(R - p_i + k_i) - \mu k_i - z\tau_1 - c \\ \text{s.t.} \quad & R - p_i + k_i \geq \xi_m(q_{bg}, \pi_{-i}, \tau_0, \tau_1) \quad \text{(IC)} \\ & sp_i + (1-s)k_i + z\tau_1 \geq I \qquad\qquad\quad \text{(IR)} \end{aligned}$$

The IR constraint binds in the optimum, and it is easy to substitute the expression for p_i into the profit function and the IC constraint. The problem becomes:

$$\begin{aligned} \max_{p_i,k_i} \quad & sR - I - (\mu - 1)k_i - c \\ \text{s.t.} \quad & sR - I + k_i + z\tau_1 \geq s\xi_m(q_{bg}, \pi_{-i}, \tau_0, \tau_1) \quad \text{(IC)} \end{aligned}$$

The optimal solution minimizes the amount of collateral k_i to be posted in order to satisfy the IC constraint. The first-best solution has $k_i = 0$, and it is implementable if and only if:

$$I \leq s[R - \xi(q_{bg}, \tau_0, \tau_1)] + z\tau_1 = I^* - s\xi_m(q_{bg}, \pi_{-i}, \tau_0, \tau_1) + c + z\tau_1 := I_c(q_{bg}, \pi_{-i}, \tau_0, \tau_1) \quad (2.15)$$

If $I > I_c(q_{bg}, \pi_{-i}, \tau_0, \tau_1)$, collateral has to be used to satisfy the IC constraint. As k_i decreases utility at a rate $\mu - 1$, the optimal choice is to set it as low as possible, i.e. to satisfy the IC constraint with equality:

$$k_i = \max\left\{0, s\xi_m(q_{bg}, \pi_{-i}, \tau_0, \tau_1) - z\tau_1 - (sR - I)\right\}$$

Substituting this quantity into the profit function, the investment is undertaken only if:

$$I \leq I^* - \frac{\mu - 1}{\mu}\left(s\xi_m(q_{bg}, \pi_{-i}, \tau_0, \tau_1) - c - z\tau_1\right) := I_k(q_{bg}, \pi_{-i}, \tau_0, \tau_1) \qquad (2.16)$$

Using equations (2.15) and (2.16), Proposition 2.5.1 characterizes the symmetric subgame perfect equilibrium of the game. The proof is analogous to the one of Proposition 2.4.1, and I omit it in the interest of space.

Proposition 2.5.1 *Let $N \to +\infty$. The only symmetric subgame perfect equilibrium is:*

- *The first-best contract $p = \frac{I - z\tau_1}{s}$, $k = 0$ for $I \leq \min\{I^*, I_c^*(q_{bg}, \tau_0, \tau_1)\}$, where $I_c^*(q_{bg}, \tau_0, \tau_1)$ is the unique solution to*

$$R - \frac{I - z\tau_1}{s} = \xi_m\left(q_{bg}, R - \frac{I - z\tau_1}{s}, \tau_0, \tau_1\right)$$

- *If* $I_c^*(q_{bg}, \tau_0, \tau_1) < I^*$, *and the investment is undertaken, dealers' final payoff* $\pi_m(\tau_0, \tau_1)$ *is the solution to* $\pi = \xi_m(q_{bg}, \pi, \tau_0, \tau_1)$. *The second-best contract is:*

$$p = \frac{I-(1-s)k - z\tau_1}{s} \qquad k = s\pi_m(\tau_0,\tau_1) - z\tau_1 - (sR - I)$$

It is profitable only if $I \in \left(I_c^*(q_{bg}, \tau_0, \tau_1), I_k(q_{bg}, \pi_m(\tau_0, \tau_1), \tau_0, \tau_1)\right]$.

- *No investment is undertaken for* $I \in \left(I_k(q_{bg}, \pi_m(\tau_0, \tau_1), \tau_0, \tau_1), I^*\right]$.

The equilibrium shares similar features with the one in the baseline model. The differences concern the determinants of the incentive compatibility constraint $\xi_m(q_{bg}, \pi, \tau_0, \tau_1)$, relative to the one with two dealers (i.e. $\xi_\tau(\pi, \tau_0, \tau_1)$). As explained before, in this model the inter-dealer market provides incentive through two channels: the threat of ostracism, and the higher incentives to avoid a match with a risky dealer. The relative importance depends on q_{bg}, and, in the remainder of this section, I analyze how the maximum investment levels change with correlation, as captured by q_{bg}. I consider separately the cases with and without CCP to better point out the effects of τ_0 and τ_1.

Equilibrium without CCP

Without a CCP, both $I_c^*(q_{bg}, \pi_{-i}, 0, 0)$ and $I_k(q_{bg}, \pi_{-i}, 0, 0)$ only depend on:

$$\xi(q_{bg}, \pi_{-i}, 0, 0) = \frac{c - q_{bg}\left[1 - G\left(\frac{L}{\gamma\pi_{-i}}\right)\right]L}{m(1-\alpha) + \gamma m(1-m)(1-\alpha-q_{bg})}$$

The correlation among dealers' shocks—which affects q_{bg}—affects the incentive compatibility constraint. Its effect depends on which mechanism makes shirking more costly: (i) matching with a safe dealer who refuses to hedge; or (ii) matching with a risky counter-party, and be exposed to greater contagion risk. An increase in q_{bg} strengthens the first channel, and weakens the second. The overall effect depends on which economic mechanism imposes the higher loss, and this trade-off is captured by the magnitude of $\mathscr{C}(\pi) = \frac{c}{\left[1-G\left(\frac{L}{\gamma\pi}\right)\right]L}$. The ratio $\mathscr{C}(\pi)$ measures the relative cost of exercising effort—and hedging for sure—against the expected cost for a shirking dealer of being matched with a safe counter-party who refuses to trade. Intuitively, a higher $\mathscr{C}(\pi)$ measures a lower expected cost from being ostracized.[17]

[17] The denominator of $\mathscr{C}(\pi)$ is $\left[1-G\left(\frac{L}{\gamma\pi}\right)\right]L$, and a lower expected cost from the threat of ostracism may come from a lower L, or from a lower probability that a safe dealer refuses to hedge.

Proposition 2.5.2 *Let π be dealers' equilibrium payoff. The incentive compatible investment levels are:*

- *Equal to I^* if*

$$\mathscr{C}\left(R-\frac{I}{s}\right) \leq \frac{\{1-\alpha m[1+\gamma(1-m)]\}\, q_{bg}}{(1-m)[1-\gamma m(1-q_{bg})]} := \eta(q_{bg})$$

  with $\eta(q_{bg})$ increasing and $\eta(0)=0$.

- *Lower than I^* if $\mathscr{C}(\pi) > \eta(q_{bg})$. The thresholds $I_c^*(q_{bg},\pi,0,0)$ and $I_k(q_{bg},\pi,0,0)$:*
  - *increase in q_{bg} if $\mathscr{C}(\pi) \in \left(\eta(q_{bg}), \frac{(1-\alpha)[1+\gamma(1-m)]}{\gamma(1-m)}\right]$.*
  - *decrease in q_{bg} if $\mathscr{C}(\pi) > \frac{(1-\alpha)[1+\gamma(1-m)]}{\gamma(1-m)}$.*

An increase in q_{bg} increases the chances to match with a safe dealer after a deviation.[18] If the expected cost of ostracism is high—i.e. $\mathscr{C}(\pi) < \frac{(1-\alpha)[1+\gamma(1-m)]}{\gamma(1-m)}$—an increase in q_{bg} relaxes the incentive compatibility constraint: it is more likely not to hedge after a deviation. In contrast, if being exposed to the contagion risk of a risky counter-party is more costly—i.e. $\mathscr{C}(\pi) > \frac{(1-\alpha)[1+\gamma(1-m)]}{\gamma(1-m)}$—a higher q_{bg} worsens the agency problem; after a deviation, the chances of matching with a risky dealer are lower.

These results can be alternatively stated as follows. If the costs of exclusion from the inter-dealer market are high, the incentives to undertake a safe business conduct are stronger—hence less collateral is necessary—in economic environments in which dealers' credit shocks depend on common macroeconomic variables. In contrast, if dealers fear more the risk of contagion, they have better incentives to be safe, and match with other safe dealers, when dealers' credit shocks are less correlated, and more dependent on idiosyncratic factors.

Equilibrium with optimal CCP design

I now turn to consider what is the potential effect of a CCP. First, I establish what is the optimal (τ_0, τ_1).

Proposition 2.5.3 *The optimal CCP penalties are $\tau_0 = 0,\ \tau_1 = \pi_{m\tau}$, where $\pi_{m\tau}$ is the solution to $\pi = \xi_m(q_{bg}, \pi, 0, \pi)$.*

The optimal loss mutualization scheme for a CCP is to impose the highest possible penalty only if a dealer has previously hedged with a defaulting member. As in the

[18] To better grasp this intuition, consider the extreme case of perfect correlation: if F_i deviates, there is $1-\alpha$ probability that all other $N-1$ dealers are safe at $t=2$.

baseline model, the penalty τ_1 enhances the incentives, for a safe dealer, to refuse a hedge from a risky counter-party. If F_i accepts a hedge, he receives a zero payoff not only when there is contagion (with prob. γ), but also when he survives, since the penalty τ_1 reduces to zero his final payoff π. However, with multiple dealers, the penalty τ_1 also increases the endogenous peer selection channel. Risky dealers are more likely to hedge with other risky dealers, and, unless they default, they are more likely to pay the penalty τ_1, as a risky hedging partner is more likely to default. Dealers want to minimize the chances of trading with a risky dealer; exercising effort increases the probability to be safe, and to trade with safe counter-parties.

Proposition 2.5.4 describes how the correlation in dealers' shocks influnces the incentive compatibility constraint and, in turn, the maximum investment level.

Proposition 2.5.4 *Let π be the equilibrium payoff with $\tau_0 = 0$ and $\tau_1 = \pi$. The maximum incentive compatible investment level is:*

- *Equal to I^* if*

$$\mathscr{C}\left(\frac{\pi}{\gamma}\right) \leq \eta(q_{bg}) + \frac{m[1-\alpha-q_{bg}(1-\alpha m)]}{[1-\gamma m(1-q_{bg})]\left[1-G\left(\frac{L}{\pi}\right)\right]L}(1-\gamma)\pi := \eta^{CCP}(q_{bg},\pi)$$

- *Lower than I^* if $\mathscr{C}\left(\frac{\pi}{\gamma}\right) > \eta^{CCP}(q_{bg},\pi)$. The thresholds $I_c^*(q_{bg},\pi,0,\pi)$ and $I_k(q_{bg},\pi,0,\pi)$:*
  - *increase in q_{bg} if $\mathscr{C}\left(\frac{\pi}{\gamma}\right) \in \left(\eta^{CCP}(q_{bg},\pi), \frac{(1-\alpha)[1+\gamma(1-m)]}{\gamma(1-m)} - \frac{m(1-\gamma)}{\gamma}\frac{(1-\alpha)\pi}{[1-G\left(\frac{L}{\pi}\right)]L}\right]$, when this interval exists.*
  - *decrease in q_{bg} if $\mathscr{C}\left(\frac{\pi}{\gamma}\right) > \max\left\{\eta^{CCP}(q_{bg},\pi), \frac{(1-\alpha)[1+\gamma(1-m)]}{\gamma(1-m)} - \frac{m(1-\gamma)}{\gamma}\frac{(1-\alpha)\pi}{[1-G\left(\frac{L}{\pi}\right)]L}\right\}$.*

The optimal CCP design improves cross-monitoring incentives. First, for a safe dealer, it is more costly to hedge with a risky counter-party; second, it is more convenient to be safe, because hedging with a risky counter-party raises the chances to pay τ_1. The first effect implies $\mathscr{C}\left(\frac{\pi}{\gamma}\right) > \mathscr{C}(\pi)$, for a given equilibrium payoff π.[19] The second effect results in a lower threshold for $\mathscr{C}(\cdot)$, above which a higher q_{bg} worsens the agency problem. Unless a first-best contract is implementable—i.e. $\mathscr{C}\left(\frac{\pi}{\gamma}\right) \leq \eta^{CCP}(q_{bg},\pi)$—the threshold is reduced by the positive term $\frac{m(1-\gamma)}{\gamma}\frac{(1-\alpha)\pi}{[1-G\left(\frac{L}{\pi}\right)]L}$. If q_{bg} increases, there is no unambiguous comparison between the equilibrium with and without CCP. Both cross-monitoring channels get reinforced, and the relative improvement—which determines the effect of q_{bg}—depends on the parameters of the game.

[19] In general the final payoff π is different between an equilibrium with or without CCP. Therefore, this comparison is not between the two equilibria.

2.6 Discussion

The results in sections 2.4 and 2.5 provide a simple framework to discuss how the organization of the inter-dealer market may, or may not, provide incentives to exploit the superior information on credit risk held by dealers. More effective peer monitoring helps to reduce the use of collateral as an incentive device.

My model points out two distinct economic mechanisms that make the inter-dealer market act as a disciplining device. First, it makes more difficult to hedge for a risky dealer, especially if he faces safe counter-parties. Second, there is endogenous peer selection, and a risky dealer is more likely to hedge with another risky dealer. In turn, a risky trading partner imposes a higher contagion risk, further reducing expected payoffs. Which mechanism is more effective depends on several variables: the number of dealers, the probability of an involuntary credit risk shock (α), and the correlation among dealers' negative shocks (q_{bg}).

Hedging ostracism is more effective when: (i) there is a small number of dealers; (ii) after exerting effort, the probability of a negative shock is small; and (iii) involuntary credit risk shocks among dealers are strongly correlated. In contrast, the incentives to be safe, and select another safe dealer, are stronger when the inter-dealer market has a large number of participants, and there is a low correlation of dealers' credit risk shocks. Under these circumstances, it is more likely that some dealers in the market are risky, and it is unlikely to have a match between a safe and a risky dealer.

Both incentive mechanisms are reinforced by the presence of a well-designed CCP. In this paper, I narrowly interpret a CCP as a financial institution[20] that impose penalties on surviving dealers upon the default of a trading partner. The optimal design is to impose a penalty on the surviving dealer only if he has previously hedged with a defaulting dealer, while it should not impose penalties if no hedge occurred. The intuition is that this penalty structure increases the cost of a risky hedge beyond the risk of contagion. Without a CCP, peer monitoring is more effective when contagion risk γ is large, but this quantity becomes irrelevant—in terms of incentive compatibility—with a CCP, since, upon the default of a dealer, it is always possible to expropriate the hedging partner payoff through a sufficiently large penalty. It is crucial that the penalty must be paid only if there was a previous hedge, otherwise it has an opposite effect on incentives, and the equilibrium would be worse than without a CCP in place.

My results have implications for the design of a CCP default waterfall. Distributing losses without any relation to previous trading patterns *harms* peer monitoring incentives. To preserve good incentives, higher initial collateral margins may be necessary.

[20] Alternatively, a CCP may be interpreted as an ex ante contract among all dealer to make them to mutually responsible of members' defaults.

According to my model predictions, most CCP default waterfalls discourage peer monitoring incentives, because—after an initial contribution out of the CCP capital—losses are distributed among members with no reference to the previous trading relationships of the defaulter. My model suggests that losses should be *first* payed by the trading partners of a defaulting member. This can be done either with additional cash contributions, or with a variation margin haircut. As a result, if peer monitoring could be a relevant tool for financial stability, a CCP default waterfall should, at least partially, *reverse the hierarchical order* of the default waterfall. Moreover, a partial haircut may reduce the risks of a CCP insolvency, as it does not require to call for additional margins, which may be difficult to raise during a market turmoil episode.

My analysis provides some sharp results on the optimal default waterfall, but it is important not to forget its caveats. First, my modelling of the inter-dealer market is extremely stylized. I exclude multiple hedging relationships, and I abstract from an endogenous determination of contagion risk in inter-dealer transactions. A more realistic modelling of inter-dealer trade would introduce a collateral decision also with respect to other dealers. This would be a first step to endogenize γ. Second, I focus on the role of peer monitoring as an incentive device, excluding any risk sharing consideration. The latter is obviously a very important concern, as CCPs have been mainly introduced for this purpose.[21] Introducing an insurance motive in the model may lead to have some default contributions from every CCP member, and not just by the defaulter's trading partners. If haircutting were a too expensive option to pursue, all CCP members would be collectively more capable of covering losses. Nonetheless, my results warn about the potential weaknesses of default waterfalls. Loss mutualization schemes that completely disregard the endogenous trade patterns among CCP members—and the valuable information they might convey—could result in higher initial margins.

2.7 Conclusion

In this paper I study peer monitoring incentives in the context of CCPs. An optimal design should take into account dealers' superior information, and reduce the use of costly commitment devices such as initial margins. Inter-dealer markets introduce two main disciplining mechanisms. First, becoming risky may lead to be excluded by other (safe) dealers, imposing additional cost from holding an unhedged position. Second, there is endogenous peer selection, and a hedge with a safe dealer reduces contagion risk. I analyze when each of these two mechanisms dominates. The threat of ostracism

[21] Although it is not clear why CCPs provide a better credit risk insurance compared to the CDS market.

is more effective when dealers' credit shocks are highly correlated, for instance as a result of a macroeconomic shock, while avoiding a hedge with a risky dealer is more important when shocks are predominantly idiosyncratic.

Both disciplining mechanisms are more effective in the presence of a CCP. Importantly, to maximize peer monitoring incentives, a CCP should impose a high enough penalty only to the hedging partners of a defaulting dealer. If all CCP members pay a similar share of losses, peer monitoring incentives would be reduced, and higher initial margins would be necessary to prevent risk-shifting. In this respect, the most common design for CCP default waterfalls seem to harm peer monitoring incentives. As a result, it could be convenient to partially reverse the order of the default waterfall. A partial haircut of contracts 'in the money' should precede additional default fund contributions by CCP members.

In future research, the model could be extended in several directions. First, dealers may endogenously choose to acquire information on other dealers' credit risk. Second, a more realistic modelling of the interactions in the inter-dealer market should lead to an endogenous determination of contagion risk.

2.8 Appendix

Proof of Proposition 2.4.1.

Step 1. F_i optimal contract is $p_i = I, k_i = 0$ if and only if $p_{-i} = I, k_{-i} = 0$.

Notice that equation (2.6) has a unique solution for $I \geq 0$. The RHS is strictly decreasing in I while the LHS is weakly increasing in I. For $I = 0$ the RHS is larger than the LHS as $R > cm$; for $I = R$ the RHS is smaller than the LHS as $c > 0$. The unique solution I_c^* falls in the relevant interval $(0, R-c]$ only if the RHS at $I = R-c$ is smaller than the LHS, i.e. if the condition stated in point 1. holds.

It easy to realize a symmetric equilibrium exist for the first-best contract $(p = I, k = 0)$ when $I \leq I_c(R-I)$. If $I_c(R-I) > I^*$ then both dealers always implement the first-best contract. If $I_c(R-I) < I^*$ then this symmetric first-best equilibrium is not implementable in $(I_c(R-I), I^*]$.

I next show there is no asymmetric equilibrium in which F_i offers the first-best contract while F_{-i} use collateral. Suppose per contra such an equilibrium exist. Then $\pi_i = R - I$ while F_{-i} payoff satisfies the IC constraint with equality, i.e. $\pi_{-i} = \frac{c - \left(1 - G\left(\frac{L}{\gamma(R-I)}\right)\right)L}{m}$. In turn, I has to simultaneously satisfy:

$$I \leq I_c(\pi_{-i}) \qquad I \in \left(I_c(R-I), I_c^k(R-I)\right]$$

Therefore, it must be $I_c(R-I) < I_c(\pi_{-i})$, or equivalently:

$$\frac{L}{\gamma(R-I)} > \frac{mL}{\gamma\left[c-\left[1-G\left(\frac{L}{\gamma(R-I)}\right)\right]L\right]}$$

Rearranging:

$$c > \left[1-G\left(\frac{L}{\gamma(R-I)}\right)\right]L + m(R-I)$$

However, substituting the expression for the lower bound $I_c(R-I)$ this inequality leads to a contradiction:

$$c > \left[1-G\left(\frac{L}{\gamma(R-I)}\right)\right]L + m\left[R-R+\frac{c}{m}-\frac{1-G\left(\frac{L}{\gamma(R-I)}\right)}{m}L\right] = c$$

Step 2. If F_i and F_{-i} offer second best contracts with $k_i, k_{-i} > 0$, they offer the same contract.

If both F_i and F_{-i} offer contracts with collateral their incentive constraints are binding and we have:

$$\pi_i = \frac{c-\left(1-G\left(\frac{L}{\gamma\pi_{-i}}\right)\right)L}{m} \qquad \pi_{-i} = \frac{c-\left(1-G\left(\frac{L}{\gamma\pi_i}\right)\right)L}{m}$$

Substituting π_{-i} in π_i, or vice versa, I get:

$$\pi_i = \frac{c-\left[1-G\left(\frac{mL}{\gamma\left[c-\left(1-G\left(\frac{L}{\gamma\pi_i}\right)\right)L\right]}\right)\right]}{m} \tag{2.17}$$

The LHS of equation (2.17) is increasing in π_i while the RHS is decreasing. Therefore, there exists a unique solution and it coincides for both F_i and F_{-i}. An easier implicit expression for $\pi_i = \pi_{-i} = \pi_c$ is:

$$\pi_c^* = \frac{c-\left[1-G\left(\frac{L}{\gamma\pi_c^*}\right)\right]L}{m}$$

Substituting the solution π_c^* for π_{-i} in $I_c^k(\pi_{-i})$ leads to the Proposition statement. Notice that equation (2.17) coincides with (2.6) for $\pi_c^* = R-I$. ■

Proof of Proposition 2.5.2.

By equations (2.15) and (2.16) the incentive compatibility constraint is not binding in a first-best solution—i.e. $I_c^*(q_{bg},\pi,0,0) > I^*$—if and only if:

$$c - s\xi(q_{bg},\pi,0,0) \geq 0$$

Substituting the expression for $\xi(q_{bg},\pi,0,0)$, it is immediate to get:

$$\mathscr{C}(\pi) = \frac{c}{\left[1-G\left(\frac{L}{\gamma\pi}\right)\right]} \leq \frac{1-\alpha m[1+\gamma(1-m)]}{(1-m)[1-\gamma m(1-q_{bg})]} q_{bg} = \eta(q_{bg})$$

with $\eta(q_{bg})$ increasing in q_{bg}.

If $\mathscr{C}(\pi) > \eta(q_{bg})$ the investment level are increasing (decreasing) in q_{bg} if and only if $\xi(q_{bg},\pi,0,0)$ is decreasing (increasing). The derivative of $\xi(q_{bg},\pi,0,0)$ with respect to q_{bg} is:

$$\frac{\mathrm{d}\xi}{\mathrm{d}q_{bg}} = \frac{\partial\xi}{\partial q_{bg}} + \frac{\partial\xi}{\partial\pi}\frac{\partial\pi}{\partial q_{bg}}$$

As π is the solution of $\pi = \xi(q_{bg},\pi,0,0)$, by implicit differentiation it is $\frac{\partial\pi}{\partial q_{bg}} = \frac{\partial\xi/\partial q_{bg}}{1-\partial\xi/\partial\pi}$. Substituting, the expression becomes:

$$\frac{\mathrm{d}\xi}{\mathrm{d}q_{bg}} = \frac{\partial\xi}{\partial q_{bg}}\frac{1}{1-\frac{\partial\xi}{\partial\pi}}$$

and the partial derivatives are:

$$\begin{aligned}
\frac{\partial\xi}{\partial q_{bg}} &= \frac{m\gamma(1-m)c - m(1-\alpha)[1+\gamma(1-m)]\left[1-G\left(\frac{L}{\gamma\pi}\right)\right]L}{[m(1-\alpha)+\gamma m(1-m)(1-\alpha-q_{bg})]^2} \\
\frac{\partial\xi}{\partial\pi} &= \frac{-q_{bg}g\left(\frac{L}{\gamma\pi}\right)\frac{L^2}{\gamma\pi^2}}{m(1-\alpha)+\gamma m(1-m)(1-\alpha-q_{bg})} \leq 0
\end{aligned}$$

Therefore, the sign of the total derivative $\frac{\mathrm{d}\xi}{\mathrm{d}q_{bg}}$ depends only on the sign of the partial $\frac{\partial\xi}{\partial q_{bg}}$. It is positive when $\mathscr{C}(\pi) \geq \frac{(1-\alpha)[1+\gamma(1-m)]}{\gamma(1-m)}$, and negative if the opposite inequality holds.

To complete the proof, it is sufficient to show that:

$$\eta(1-\alpha) = \frac{1-\alpha[1+\gamma(1-m)]}{(1-m)(1-\gamma\alpha m)}(1-\alpha) < \frac{(1-\alpha)[1+\gamma(1-m)]}{\gamma(1-m)}$$

since $\eta(q_{bg})$ is increasing in q_{bg}. Rearranging this expression, it is leads to $1-\gamma+\gamma(1-m) > 0$, which always holds. ■

Proof of Proposition 2.5.3.

By equations (2.15) and (2.16), the maximum investment level depends on (τ_0, τ_1) through the quantity:

$$\mathscr{J} = z\tau_1 + c - s\xi\left(q_{bg}, \pi(\tau_0, \tau_1), \tau_0\tau_1\right)$$

Suppose $I_c^*\left(q_{bg}, \pi, \tau_0, \tau_1\right) < I^*$. The total derivative of $\mathscr{J}$ with respect to τ_1 is:

$$\frac{\mathrm{d}\mathscr{J}}{\mathrm{d}\tau_1} = z - s\left[\frac{\partial\xi}{\partial\pi}\frac{\partial\pi}{\partial\tau_1} + \frac{\partial\xi}{\partial\tau_1}\right]$$

By Proposition 2.5.1, π is the solution to $\pi - \xi\left(q_{bg}, \pi(\tau_0, \tau_1), \tau_0\tau_1\right) = 0$. By implicit differentiation of this equation I get:

$$\frac{\partial\pi}{\partial\tau_1} = \frac{\frac{\partial\xi}{\partial\tau_1}}{1 - \frac{\partial\xi}{\partial\pi}}$$

Substituting in the expression for the total derivative and simplifying:

$$\frac{\mathrm{d}\mathscr{J}}{\mathrm{d}\tau_1} = z - s\frac{\frac{\partial\xi}{\partial\tau_1}}{1 - \frac{\partial\xi}{\partial\pi}}$$

This expression is positive as:

$$\begin{aligned}
\frac{\partial\xi}{\tau_1} &= -\frac{q_{bg}g\left(\frac{L}{h}\right)\frac{L^2}{h^2}(1-\gamma) + m(1-m)(1-\alpha-q_{bg})(1-\gamma)}{m(1-\alpha) + \gamma m(1-m)(1-\alpha-q_{bg})} \leq 0 \\
\frac{\partial\xi}{\partial\pi} &= -\frac{q_{bg}g\left(\frac{L}{h}\right)\frac{L^2}{h^2}\gamma}{m(1-\alpha) + \gamma m(1-m)(1-\alpha-q_{bg})} \leq 0
\end{aligned}$$

Therefore, it is optimal to set τ_1 as high as possible, i.e. $\tau_1 = \pi$.

An analogous calculation for the total differential with respect to τ_0 leads to:

$$\frac{\mathrm{d}\mathscr{J}}{\mathrm{d}\tau_0} = -s\left[\frac{\partial\xi}{\partial\pi}\frac{\partial\pi}{\partial\tau_0} + \frac{\partial\xi}{\partial\tau_0}\right] = -s\frac{\frac{\partial\xi}{\partial\tau_0}}{1 - \frac{\partial\xi}{\partial\pi}} \leq 0$$

since $\frac{\partial\xi}{\partial\tau_0} = \frac{q_{bg}g\left(\frac{L}{h}\right)\frac{L^2}{h^2}}{m(1-\alpha)+\gamma m(1-m)(1-\alpha-q_{bg})} \geq 0$ and $\frac{\partial\xi}{\partial\pi} < 0$. Therefore, $\tau_0 = 0$ is optimal. ■

Proof of Proposition 2.5.4.

The first-best investment level is always attainable if:

$$\mathscr{J} = z\pi + c - s\xi\left(q_{bg}, \pi, 0, \pi\right) \geq 0$$

Inserting the expression in equation (2.14), the threshold can be expressed as:

$$\mathscr{C}\left(\frac{\pi}{\gamma}\right)=\frac{c}{\left[1-G\left(\frac{L}{\pi}\right)\right]L]}\leq\eta(q_{bg})+\frac{m[1-\alpha-q_{bg}(1-\alpha m)]}{[1-\gamma m(1-q_{bg})]\left[1-G\left(\frac{L}{\pi}\right)\right]L}(1-\gamma)\pi=\eta^{CCP}(q_{bg},\pi)$$

When the opposite inequality holds the effect of q_{bg} can be computed through the total derivative with respect to q_{bg}, i.e.:

$$\frac{\mathrm{d}\mathscr{J}}{\mathrm{d}q_{bg}}=-s\left[\frac{\partial\xi}{\partial q_{bg}}+\frac{\partial\xi}{\partial\pi}\frac{\partial\pi}{\partial q_{bg}}\right]+z\frac{\partial\pi}{\partial q_{bg}}$$

By implicitly differentiating equation $\pi-\xi(q_{bg},\pi,0,\pi)=0$ with respect to q_{bg} I get:

$$\frac{\partial\pi}{\partial q_{bg}}=\frac{\frac{\partial\xi}{\partial q_{bg}}}{1-\frac{\partial\xi}{\partial\pi}}$$

Substituting in the total derivative:

$$\frac{\mathrm{d}\mathscr{J}}{\mathrm{d}q_{bg}}=-(s-z)\frac{\frac{\partial\xi}{\partial q_{bg}}}{1-\frac{\partial\xi}{\partial\pi}}$$

Since:

$$\begin{aligned}
s-z &= \alpha(1-m)^2+(1-\alpha)>0\\
\frac{\partial\xi}{\partial\pi} &= -\frac{q_{bg}g\left(\frac{L}{\pi}\right)\frac{L^2}{\pi^2}+m(1-m)(1-\alpha-q_{bg})(1-\gamma)}{\left[m(1-\alpha)+\gamma m(1-m)(1-\alpha-q_{bg})\right]^2}<0\\
\frac{\partial\xi}{\partial q_{bg}} &= \frac{m\gamma(1-m)c-(1-\alpha)[1+\gamma(1-m)]\left[1-G\left(\frac{L}{\pi}\right)\right]L+m^2(1-m)(1-\alpha)(1-\gamma)\pi}{m^2[1+\gamma(1-q_{bg})(1-m-\alpha)]^2}
\end{aligned}$$

Hence, the sign of $\frac{\mathrm{d}\mathscr{J}}{\mathrm{d}q_{bg}}$ is determined by $\frac{\partial\xi}{\partial q_{bg}}$. It is positive when:

$$\mathscr{C}\left(\frac{\pi}{\gamma}\right)=\frac{c}{\left[1-G\left(\frac{L}{\pi}\right)\right]L}\geq\frac{(1-\alpha)[1+\gamma(1-m)]}{\gamma(1-m)}-\frac{m(1-\gamma)}{\gamma}\frac{(1-\alpha)\pi}{\left[1-G\left(\frac{L}{\pi}\right)\right]L}$$

and negative otherwise. ■

Bibliography

ACHARYA, V., AND A. BISIN (2014): "Centralized versus Over-The-Counter Markets," *Journal of Economic Theory*, 149, 153–182.

AFFINITO, M. (2012): “Do Interbank customer relationships exist? And how did they function in the crisis? Learning from Italy,” *Journal of Banking & Finance*, 36, 3163–3184.

ANTINOLFI, G., CARAPELLA, F., AND F. CARLI (2014): “Clearing, Transparency, and colleteral,” working paper.

BIAIS, B., HEIDER, F., AND M. HOEROVA (2012): “Clearing, Counterparty Risk and Aggregate Risk,” ECB working paper series N. 1481.

BRÄUNING, F., BLASQUES, F., AND I. VAN LEVYLEND (2014): “A Dynamic Networl Model of the Unsecured Interbank Lending Market,” working paper.

CARAPELLA, F., AND D. C. MILLS (2012): “Information Insensitive Securities: the Benefits of Central Counter-parties,” working paper.

CARAPELLA, F., AND C. MONNET (2013): “Dealers’ Insurance, Market Structure and Liquidity,” working paper.

CRAIG, B., KOETTER, M., AND U. KRÜGER (2014): “Interbank lending and distress: observables, unobservables, and network structure,” Deutsche Bundesbank Discussion Paper N. 18.

DUFFIE, D., AND H. ZHU (2011): “Does a Central Clearing Counterparty Reduce Counterparty Risk?,” *Review of Asset Pricing Studies*, 1, 74–95.

ELLIOTT, D. (2013): “Central Counterparty Loss-Allocation Rules,” *Bank of England Financial Stability Paper* No. 20.

GHATAK, M. (2000): “Screening by the Company You Keep: Joint Liability Lending and the Peer Selection Effect,” *The Economic Journal*, 110, 601–631.

KOEPPL, T. V. (2013): “The Limits of Central Counterparty Clearing: Collusive Moral Hazard and Market Liquidity, ” working paper.

KOEPPL, T. V., AND C. MONNET (2012): “Central Counterparty Clearing and Systemic Risk Insurance in OTC Derivatives Markets,” working paper.

KOEPPL, T. V., AND C. MONNET (2013): “The Emergence and Future of Central Counterparties, ” working paper.

LIEDORP, F.R., MEDEMA, L., KOETTER, M., KONING, R.H, AND I. VAN LELYVELD (2010): “Peer monitoring or contagion? Interbank market exposure and bank risk,” DNB working paper N. 249.

MIGLIETTA, A., PICILLO, C., AND M. PIETRUNTI (2015): "The Impact of CCPs' margin policies on Repo Markets," working paper.

ROCHET, J.C., AND J. TIROLE (1996): "Interbank Lending and Systemic Risk," *Journal of Money, Credit and Banking*, 28, 733–762.

STIGLITZ, J. E. (1990): "Peer Monitoring and Credit Markets," *The World Bank Economic Review*, 4, 351–366.

VARIAN, H. R. (1990): "Monitoring Agents With Other Agents," *Journal of Institutional and Theoretical Economics*, 146, 153–174.

Chapter 3

Learning and Price Dynamics in Durable Goods Markets

A durable good provides a private use value to its user, and it is eventually resold in a secondary market. This paper analyzes what determines different learning and price dynamics in durable goods markets. Our model includes three main features: (i) buyers have heterogenous private use values and a common expected resale horizon; (ii) an unobservable and time-varying aggregate state determines the distribution of use values in the population; and (iii) trade takes place in markets with a limited number of buyers. Informational frictions slow down learning and affect price movements asymmetrically in high and low aggregate states. We disentangle two sources of price variability. Idiosyncratic volatility is prevalent in markets with very heterogenous use values, a long resale horizon and a small number of buyers. Aggregate volatility mirrors the sensitivity of prices to new price information, and it weights more when the resale motive dominates, i.e. for goods with short resale horizons, significant persistence of the aggregate state, and similar use values.

3.1 Introduction

Since the initial contribution of Hayek (1945), a vast literature in economic theory has been studying how the price system aggregates dispersed private information. No social planner has access to all available information, and in a market based economy prices have the fundamental of influencing decisions by consumers, firms and governments. How information is *incorporated* into asset prices is the main focus of the Rational Expectation literature, and one of the most debated topics in Finance.

Although there is a vast asset pricing literature on financial securities, less attention has been devoted to price patterns in durable goods markets. Notable example are real-estate, machineries, automotive, but also artwork, collectibles and musical instruments. These goods provide a private use value to users, but they are often resold on the market after some time. Durable goods represent a sizeable portion of household and corporate balance sheets, and as such they play a central role in the economy as consumption goods, production inputs or pledgeable collateral. Many papers focus on a specific market—especially real-estate and vehicles—and try to match a few empirical facts, either with a rather specific model, or with a slight adaptation of a workhorse asset pricing model. In the former case, the model results cannot be applied *sic et simpliciter* to other markets which share a few similarities; in the second case, models overlook some specific—but potentially relevant—market features.[1]

In this paper we broadly focus on durable goods—a sufficiently large class of assets—and we study how a few common characteristics affect learning and information aggregation. Our model does not pretend to match precise price patterns for a specific market,[2] but it rather aims to highlight a few economic mechanisms potentially relevant to all durable goods markets.

We develop a dynamic trading model with time-varying and unobservable aggregate demand conditions. Our framework explicitly considers two peculiar characteristics of durable goods. First, they trade in decentralized markets where sellers enter into private negotiations with a limited number of potential buyers. Second, they provide utility as consumption goods until re-sold to a different user at a future point in time.[3] There exist great variation within each characteristic. On the one hand, trade decentralization admits a large variety of trade protocols. On the other hand, the consumption vs. resale trade-off depends on several intrinsic characteristics of the market.

[1] An example of this dihcotomic approach is the real estate literature. Some authors use Lucas-types models and derive estimates for risk and liquidity premia, other papers set up search and matching models including a rental sector, geographic dispersion, and private use values.

[2] We do not deal with any specific price puzzle, and we actually exclude a priori the existence of risk-premia by assuming agents' risk neutrality. Our main focus is on information.

[3] Other products may share the same two features. We explcitly refer to durable goods just to focus our attention on a relevant set of markets which possess these broad characteristics.

Learning patterns depend on prices if the latter provide useful information on the underlying aggregate demand. Heterogeneity in trading protocols leads to different ways in which agents update their beliefs. These informational frictions may have different origins: the absence of an organized trading platform, legal restrictions on information disclosure, or bidders' incentives to manipulate prices. We abstract from any single source of friction, and we focus directly on the relationship between disclosed information and learning dynamics. We present two main results. First, trading games revealing *coarser* information sets leat to a *slower* learning process. Second, different trading protocols may affect beliefs *asymmetrically* between high and low aggregate demand states. In particular, when only winning bids are disclosed, beliefs tend to adjust more rapidly when the aggregate state is low.

If the trading protocol determines which information is revealed to agents, other intrinsic characteristics of the durable good influence its price sensitivity to new information. We consider three main dimensions: the expected resale horizon, the persistence of aggregate demand states, and the degree of heterogeneity in private use values. To explicitly solve the model, we assume sellers trade via second-price auctions. Thanks to an analytic solution for the bidding strategy, we obtain several comparative statics results. First, prices respond more to new information when buyers have more similar private use values. Second, a longer expected resale horizon increases the relative importance of private use values *vis-á-vis* future resale prices. Similarly, price sensitivity is larger when aggregate states are more persistent. Lastly, price volatility can be decomposed into two factors: idiosyncratic and aggregate. The former depends on the heterogeneity in buyers' use values, and it is driven by the consumption motive. The latter captures price sensitiveness to current information, and it depends on the interest in forecasting future prices.

Despite theoretical in nature, we believe our paper points out a few general ideas with a broad range of potential applications. For example, suppose a credit officer has to decide on the loan terms applied to two otherwise identical customers with different collateral goods: one has an classic car, and the other one a modern corporate car. Which car is the less risky collateral? To answer this question, it might be a good idea to understand who participates in these markets, and for which purpose. Classic cars are mostly bought for their subjective use value, and usually resold after a long time. On the contrary, buyers of corporate cars have similar use values, a fast car turnover, and they significantly care about the future resale price. Our model provides a framework to explain how these different characteristics may affect price volatility.

Overview of the model and results. We briefly sketch our model setup to discuss our results in more detail. Trade takes place through a sequence of trading rounds with N bidders. Aggregate market conditions in period t depend on the distribution of private values from which individual bidders are sampled. In particular, their per-period use value in period t come from one of two possible distribution functions F_{θ_t}, $\theta_t \in \{H, L\}$. The state of the world θ_t is never publicly revealed, and it varies overtime according to a Markov process with state persistence ρ_j, $j = H, L$. Unless θ_t realizations are independent overtime, the observable public history provides information on the likelihood of future states of the world. Private use values have a double role: (i) they measure individual benefits from enjoying the good; and (ii) they provide information on the underlying state of the world. A winning bidder resells his good at a future random time: he faces a $\alpha \leq 1$ probability to sell his good in the next period. Higher values of α denote shorter resale horizons.[4] An owner enjoys his individual use value until resale. For simplicity, losing bidders and sellers go out of the market with no future possibility to re-enter.

The aggregate state θ_t may be considered as a reduced form to capture all those elements such as fashion, business and credit cycles that affect, at a given point in time, the willingness, or possibility, to purchase the good among agents in the population. It is often difficult to directly observe this state and we assume buyers only observe previous transaction prices. For example, a real-estate buyer may collect information on past prices in a local market but he may not have (or be able to process) information on unsuccessful bids, or on the real-estate market at large.

Our setup captures a few characteristics of market demand that widely vary across durable goods. The parameter α is a reduced form to capture the expected resale horizon for the good. The state persistence parameter ρ measures how likely an aggregate state will persist in future periods; lower values of ρ_H and ρ_L denote a more volatile aggregate environment. The distributions F_θ describe a more or less dispersed distribution of private use values among agents in the economy. Finally, the number of bidders N provides a measure of market competition, but also, to a certain extent, market liquidity. Thanks to an explicit characterization of the bidding function it is possible to derive analytically some general comparative statics results, and it would be straightforward to simulate other statistical properties for specific functions F_θ.

In section 3.2 we discuss how differences in the information revealed through prices affect learning dynamics. The more information is disclosed by a trade protocol, the faster beliefs converge to the true state. In this respect, durable goods markets may exhibit a more sluggish price adjustment process relative to a centralized market.[5] The second result is less intuitive. In general, the speed of learning differs between high and low aggregate states. This asymmetry depends on the information revealed by the

[4] Alpha may be considered a measure of the likelihood to be hit by a liquidity shock that forces to sell the object.

[5] For example, in a modern stock exchange dealers' price quotes and traders' limit orders can be freely observed in real time by all market participants. Without strategic price manipulation markets disclose all private information.

trade protocol. For example, a first-price auction reveals the highest valuation among the N bidders. In this case, learning is faster in the low state because low prices are more informative in revealing the underlying aggregate demand state.[6]

In section 3.3 we assume—for reasons of analytical tractability—that the object is sold in a sequence of second-price auctions.[7] Prices are more sensitive to new information in markets in which: (i) the resale horizon is shorter ($\alpha \uparrow$); (ii) the expected demand between high and low aggregate states is larger; (iii) the current state of market demand is more likely to last longer ($\rho_j \uparrow$). Under these circumstances, sellers weight more the informational content of recent prices, as the latter are more effective in predicting future resale values. Price variability can be decomposed in two different components. The first one reflects the heterogeneity in private use values, and it has a purely idiosyncratic nature. The second type of uncertainty is over future market conditions. A decrease in the resale horizon ($\alpha \uparrow$) increases the aggregate variability component, decreasing the idiosyncratic one; thus, the overall effect is ambiguous. A increase in state persistence ($\rho_j \uparrow$) does not affect the idiosyncratic variance, and it increases the variability due to the future resale component.

Related literature. This paper is closely related to the literature on learning in asset pricing models; see Timmermann (1993) for an early reference. This strand of literature argues that learning may explain some classic asset pricing puzzles (equity premium, risk-free rate and excess-volatility). Weitzman (2007) considers a Bayesian framework with risk-averse preferences, while Ju and Miao (2012) introduce ambiguity aversion as an additional explanatory factor. Compared to this literature, we focus more on the different microeconomic determinants of price dynamics, and we abstract from any discussion on risk-premia by assuming risk neutral preferences.

This paper is also related to the literature on auctions with resale. A small number of papers study this topic in a two periods setting. Gupta and Lebrun (1999) consider a setup in which private values are publicly revealed in the second period. Haile (2001, 2003) study the revenue performance of different auction formats in a symmetric environment. In his model, bidders' initial types come from the same initial distribution but they are not publicly announced in the second period. Within a similar symmetric environment, Zheng (2002) and Lebrun (2012) provide conditions to obtain the optimal auction outcomes first derived in Myerson (1981).[8] Differently from this literature, we do not assume that the same set of bidders re-trades in future periods. The latter case

[6] An analogous, but opposite, logic would hold if the lowest valuation among N bidders were revealed.

[7] In a second-price auction, players do not engage in strategic price manipulation: the optimal bidding strategy truthfully reveals private signals. In other auction formats—such as the first-price auction—manipulative incentives may arise, and the equilibrium analysis becomes analytically intractable.

[8] Lebrun (2012) extend the results in Zheng (2002) for the symmetric environment, and it provides results for a specific class of asymmetric environments. Garratt and Troger (2006), Halafir and Krishna (2008) and Virag (2013) provide additional results for asymmetric environments.

is important for industries in which market players rarely change overtime,[9] and manipulative incentives may arise when the same good is re-traded among the same set of bidders. However, in many durable goods markets this type of strategic interaction seems less relevant. For example, in the real-estate market, buyers and sellers often do not have any previous information on the identity of their counterpart.[10]

The next section discusses some general results on public learning. Section 3 presents the dynamic auction model, and it provides comparative statics. Section 4 concludes. All proofs are in the Appendix.

3.2 Information revelation and learning

3.2.1 Model setup

We consider a sequential market for a durable object. Time is discrete $t \in \{0,1,2,\ldots\}$. In each period t, there is an underlying state $\theta_t \in \{H, L\}$. The stochastic process $\{\theta_t\}_{t=0}^{\infty}$ is a homogenous Markov process with transition matrix:

$$P = \begin{bmatrix} \rho_H & 1-\rho_H \\ 1-\rho_L & \rho_L \end{bmatrix}$$

with $0 \leq \rho_H, \rho_L \leq 1$. The prior on θ_t is denoted by $\pi_t = (\pi_t, 1-\pi_t)$ with $\pi_t \equiv \mathbb{P}(\theta_t = H)$.

There is a population of infinitely many agents interested in the object. When an object is offered on sale, $N \geq 2$ agents are randomly drawn from the population to enter the market. Each buyer attaches a private use value to the object. The private values generated in each period t, $\{v_{it}\}_{i=1}^{N}$, are i.i.d. distributed according to a cumulative density function (cdf) F_{θ_t} across the N agents. The realizations of $\{\theta_t\}_{t=0}^{\infty}$ are not known to the agents, but both P and π_0 are common knowledge.

Both F_H and F_L are continuously differentiable on the common support $[0,1]$. Moreover, the corresponding probability density function (pdf) f_H and f_L are strictly positive everywhere on $[0,1]$, and satisfy the monotone likelihood ratio (MLR) property: $\frac{f_H(\cdot)}{f_L(\cdot)}$ is strictly monotone on $[0,1]$.

[9] For example, government concessions in telecommunications, oil, electricity.

[10] Similar to our model, in some Australian cities a significant portion of home sales take place via auctions; see www.bloomberg.com/news/2013-04-23/australia-turns-to-auctions-as-housing-revives-mortgages.html

3.2.2 Public beliefs dynamics

The information revealed in a trading round depends on the trade protocol. For example, there is a substantial difference between auctions and centralized exchanges, but there are also significant differences among auction formats. In this section, we abstract from a specific trade protocol, and we directly consider the information revealed after a trading round.[11]

Consider a vector $X_k = \{v_{1,k}, ..., v_{N,k}\}$ of private signals dispersed among N traders. We assume the trade protocol leads to publicly observe a statistic $T(X_k)$. $T(\cdot)$ is assumed not to depend on the previous history of the game, hence it is invariant in all periods k. In other words, the statistic T captures which information in the vector of private valuations X_k possessed by the N bidders in period k is publicly revealed after trade.[12] It is an equilibrium object because it depends on the trade protocol and players' strategies. It implicitly incorporates both informational constraints, due to the market organization, and informational frictions, due to players' strategic behaviour.

We explore two different issues related to learning. First, we provide a sufficient condition that ranks which statistic leads to a faster public belief convergence towards the true state. Second, we analyze whether a particular T leads to a more rapid price adjustment in one of the two states of the world. For these purposes, we restrict attention to the full persistence case $\rho_H = \rho_L = 1$.

Consider a probability space $\langle \mathbb{R}^N, \mathscr{B}, \mu \rangle$ endowed with the standard Borel σ-algebra and Lebesgue probability measure. Let a measurable function $T_i : \mathbb{R}^N \to \mathbb{R}^M$, $M \leq N$ be an observable statistic of the underlying $X_k = \{v_{1,k}, ...v_{N,k}\}$ and let $\sigma(T_i)$ be the σ-algebra generated by T_i. I denote with S_X the support of X.

Definition 3.2.1 *T_j is coarser than T_i if $\sigma(T_j) \subset \sigma(T_i)$ and $\exists A \in \sigma(T_i)$ s.t. $A \notin \sigma(T_j)$ and $\mu(A) > 0$.*

For a statistic T let S_T denote its support and $\mathscr{C}_T(A) \equiv \bigcup_{y \in A} \left\{X \in \mathbb{R}^N : T(X) = y\right\}$ the set of counter images of $A \subseteq S_T$.

We use this general notation to express public belief dynamics under different trade protocols. Let $f_\theta^T(y)$ be the probability density function of statistic T under state $\theta \in \{H, L\}$. Formally,

$$f_\theta^T(y) = \int_{\mathscr{C}_T(y)} f_\theta^X(x)\, \mathrm{d}\mu(x)$$

It is easier to describe the evolution of public beliefs with the log-likelihood ratio:

[11] In section 3.3, we solve a specific model where agents participate in second-price auctions. In this section, we adopt a more general approach to point out a few general properties of learning dynamics.

[12] For example, in section 3.3.2 the second-highest price is publicly revealed and in equilibrium buyers infer the corresponding private value, hence $T(X_k) = v_k^{(2,N)}$.

$$l_{k+1}(l_k, y) = \ln \frac{\pi_{\kappa+1}}{1-\pi_{\kappa+1}} = \ln \frac{\pi_\kappa}{1-\pi_\kappa} \frac{f_H^T(y_k)}{f_L^T(y_k)} = l_k + \ln \frac{f_H^T(y_k)}{f_L^T(y_k)} \tag{3.1}$$

where $y_k = T(x_k)$ is the value of statistic T when the vector of private use values for the N bidders in period k is $x_k \in \mathbb{R}^N$. To stress that the log-likelihood l_k depends on T, we add a superscript T. Let $\Delta l_{k+1}^T(y_k) = l_{k+1}^T(l_k^T, y_k) - l_k^T = \ln \frac{f_H^T(y_k)}{f_L^T(y_k)}$ denote the change in the log-likelihood ratio from period k to $k+1$ under statistic T. Assume there exists $M > 0$ such that $|\Delta l^T(x)| < M$ for every $x \in S_X$.

For $q \geq 1$ equation (3.1) generalizes into:

$$l_{k+q}^T = l_k^T + \sum_{m=1}^{q} \Delta l_{k+m}^T(y_{k+m-1})$$

Taking the expected value:

$$\mathbb{E}_{k,\theta}^X[l_{k+q}^T] = l_k^T + \sum_{m=1}^{q} \mathbb{E}_{k,\theta}^X\left[\Delta l_{k+m}^T(T(x_{k+m-1}))\right] = l_k^T + q\mathbb{E}_\theta^X\left[\Delta l^T(T(x))\right]$$

The last equation exploits the fact that—conditional on θ—samples are i.i.d. in all periods. In the remainder of the paper, we simply use $\mathbb{E}_\theta\left[\Delta l^T\right]$ rather than $\mathbb{E}_\theta^X\left[\Delta l^T(T(x))\right]$.

Beliefs converge to the true state as $\mathbb{E}_H\left[\Delta l^T\right] > 0 > \mathbb{E}_L\left[\Delta l^T\right]$.[13] Moreover, for two different statistics T_1 and T_2, public beliefs are expected to converge more rapidly to the true state θ under statistic T_1 if:

$$\left|\mathbb{E}_\theta\left[\Delta l^{T_1}\right]\right| > \left|\mathbb{E}_\theta\left[\Delta l^{T_2}\right]\right| \tag{3.2}$$

The next lemma provides an intuitive but still insightful result. If two statistics can be ranked according to Definition 3.2.1, it is possible to conclude that convergence is slower for the coarser one.

Lemma 3.2.1 *If T_2 is coarser than T_1 then equation (3.2) holds.*

Although the result in Lemma 3.2.1 is not surprising, it highlights an important property of markets in which information is only partially revealed. More severe informational frictions lead to more sluggish trade dynamics, and, possibly, a slower price adjustment.

A less intuitive result is that trade protocols may create differences—between high and low states—in the speed of convergence of public beliefs. In turn, more rapid learning is likely to be positively correlated with a more rapid price adjustment.

[13] It follows from a simple application of Gibbs' inequality.

Lemma 3.2.2 *Consider a statistic $T(\cdot)$, and let $\pi_0 = 1/2$. Define:*

$$\tau_H := \inf\{k \geq 0 : \pi_k \geq 1-\varepsilon\} \qquad \tau_L := \inf\{k \geq 0 : \pi_k \leq \varepsilon\}$$

Then:

$$\lim_{\varepsilon \to 0} \frac{\mathbb{E}_H[\tau_H]}{\mathbb{E}_L[\tau_L]} \geq \left|\frac{\mathbb{E}_L[\Delta l^T]}{\mathbb{E}_H[\Delta l^T]}\right| > 1 \quad \text{if } \mathbb{E}_H[\Delta l^T] + \mathbb{E}_L[\Delta l^T] < 0$$
$$\lim_{\varepsilon \to 0} \frac{\mathbb{E}_L[\tau_L]}{\mathbb{E}_H[\tau_H]} \geq \left|\frac{\mathbb{E}_H[\Delta l^T]}{\mathbb{E}_L[\Delta l^T]}\right| > 1 \quad \text{if } \mathbb{E}_H[\Delta l^T] + \mathbb{E}_L[\Delta l^T] > 0$$

Lemma 3.2.2 points out a learning story based on the nature of the information revealed in previous trading rounds. Compared to the 'rockets and feathers' story, our mechanism is likely to run the opposite way. If a trade protocol only reveals winning bids, a more rapid adjustment should be observed downward. We discuss the intuition in the context of an example.

Example Lemma 3.2.2. Consider pdfs $f_H(x) = 2x$ and $f_L(x) = 2(1-x)$. Suppose the trade protocol reveals, in equilibrium, the j-th order statistic out of N bidders. The next table summarizes the numerical values of the condition in Lemma 3.2.2:[14]

$N \mid j$	1	2	3	4	5	6	7	8	9	10
1	0									
2	−0.09	0.09								
3	−0.22	0	0.22							
4	−0.37	−0.13	0.13	0.37						
5	−0.52	−0.28	0	0.28	0.52					
6	−0.68	−0.44	−0.14	0.14	0.44	0.68				
7	−0.83	−0.60	−0.30	0	0.30	0.60	0.83			
8	−0.99	−0.77	−0.46	−0.15	0.15	0.46	0.77	0.99		
9	−1.14	−0.94	−0.63	−0.31	0	0.31	0.63	0.94	1.14	
10	−1.29	−1.12	−0.80	−0.48	−0.16	0.16	0.48	0.80	1.12	1.29

The value is negative (positive) when beliefs move more rapidly toward state L (H), as the expecting hitting time under state L (H) is shorter. The table shows that:

1. If $j < \frac{N+1}{2}$ convergence is faster towards state L.
2. If $j > \frac{N+1}{2}$ convergence is faster towards state H.
3. If $j = \frac{N+1}{2}$, (N odd), there is no difference.

The table captures an intuitive result. If the trade protocol reveals a sequence of higher order statistic ($j < \frac{N+1}{2}$), low-value observations are more informative than

[14] As explicit integrals cannot be obtained we compute integrals numerically.

high-value ones, and learning in more rapid in state L.[15] Fixing j and increasing the sample size N, there is more and more asymmetry toward state L.[16] Increasing the sample size N, the j-th order statistic is relatively 'higher', low-value observations become more informative, and there is a greater asymmetry towards state L. For example, if only the winning bid is revealed, low demand states are learnt more rapidly in a large market.

In conclusion, prices are not equally informative on both aggregate states. This phenomenon depends on the original distribution functions F_θ, but also on the trading protocol and the number of market participants.

3.3 Dynamic auction model

3.3.1 Trading protocol

Consider the model setup in section 3.2.1. Now assume agents trade in second-price auctions according to the following protocol.

1. Consider a stochastic sequence $\{t_k\}_{k=0}^{\infty}$, with t_0 normalized to 0. At each t_k, N new agents enter the market, and participate in a sealed-bid second- price auction. The winner of the auction at t_k is the seller in the next available auction at t_{k+1}.

2. The waiting time between t_{k+1} and t_k is a random variable, which is i.i.d distributed across k according to a geometric distribution with parameter $\alpha \in (0,1]$. That is,

$$\mathbb{P}(\triangle_k \equiv t_{k+1} - t_k = x) = \alpha(1-\alpha)^{x-1}, \forall x \in \mathbb{N}^+, \forall k \in \mathbb{N}.$$

   Due to the i.i.d. feature of the waiting time, we simply call the auction at t_k as "auction k". We also label each bidder in auction k by ik, $i \in \{1,2,\dots,N\}$. Denote his private value and bid as $w_{ik} \equiv v_{i,t_k}$ and b_{ik}, respectively.

3. The winner of auction k resells the object at the next available auction $k+1$. The revenues from resale are discounted at rate δ per period. Meanwhile, he enjoys his private value of the object, w_{ik}, in every period before auction $k+1$, and he discounts his utility at rate δ per period.

[15] On the contrary, if a sequence of lower order statistic ($j > \frac{N+1}{2}$) is revealed, high-value observations are more informative than low-value ones, and learning is more rapid in state H.

[16] Alternatively, there is less asymmetry towards H as values get smaller downwards in each column.

4. The trading price $p_k \equiv b_k^{(2)}$, the second highest bid in auction k, is publicly observed by the whole population before the next auction starts. There is no information generated between two adjacent auctions, other than the realization of the waiting time in between. Hence, the information set for each bidder ik is $\mathscr{I}_{ik} = \{w_{ik}, \{p_\tau\}_{\tau<k}, \{\Delta_\tau\}_{\tau<k}\}$.

3.3.2 Equilibrium characterization

Let $\mathbf{b} \equiv \{b_{ik}\}_{i\leq N, k\in\mathbb{N}}$ denote the action profile of all market entrants and every bidder ik's payoff is given by:

$$u_{ik}(\mathbf{b}; w_{ik}) = 1_{\{b_{ik}=b_k^{(1)}\}} \cdot \{-b_k^{(2)} + \sum_{s=0}^{\triangle_{k-1}} \delta^s w_{ik} + \delta^{\Delta_k} b_{k+1}^{(2)}\}$$

Denote the public belief about the underlying state θ before auction k by $\pi_k = (\pi_k, 1-\pi_k)$ with $\pi_k \equiv \mathbb{P}(\theta_{t_k} = H | \{p_\tau\}_{\tau<k})$.

We consider a Perfect Bayesian Equilibrium with *symmetric, time-invariant and monotone* strategies. In turn, every bidder ik can restrict his attention to the information set $\{v_{ik}, \pi_k\}$.

Definition 3.3.1 *A pure strategy profile* $\mathbf{b}^* \equiv \{b_{ik}^*(\mathscr{I}_{ik})\}_{i\leq N, k\in\mathbb{N}}$ *is a perfect Bayesian equilibrium with symmetric, time-invariant and monotone strategies if:*

1. $b_{ik}^*(\mathscr{I}_{ik}) = b^*(w_{ik}; \pi_k)$, $\forall i, t$;

2. $\frac{\partial b^*(w_{ik};\pi_k)}{\partial w_{ik}} > 0$, $\forall w_{ik} \in [0,1]$;

3. $b^*(w_{ik}; \pi_k) = \arg\max_b \mathbb{E}[u_{ik}(b, b_{-ik}^*; w_{ik}) | w_{ik}, \pi_k]$, $\forall i, k$;

4. $\pi_{k+1} = \left(\frac{\pi_k g_H(b^{*-1}(p_k;\pi_k))}{\pi_k g_H(b^{*-1}(p_k;\pi_k)) + (1-\pi_k) g_L(b^{*-1}(p_k;\pi_k))}, \frac{(1-\pi_k) g_L(b^{*-1}(p_k;\pi_k))}{\pi_k g_H(b^{*-1}(p_k;\pi_k)) + (1-\pi_k) g_L(b^{*-1}(p_k;\pi_k))} \right) P^{\Delta_k}$

   where $g_\theta(\cdot)$ *is the pdf of the* 2^{nd} *order statistic among N i.i.d random variables distributed according to* F_θ, $\forall \theta \in \{H, L\}$.

The next Proposition provides an explicit characterization of the bidding strategy.

Proposition 3.3.1 *Let* $\rho_H + \rho_L \geq 1$. *There is a unique perfect Bayesian equilibrium with symmetric, time-invariant and monotone strategies:*

$$b^*(w_{ik}; \pi_k) = \frac{1}{1-\delta+\alpha\delta} \left\{ w_{ik} + \frac{\alpha\delta}{1-\delta} \left[c_L + \frac{1-\rho_L}{1-\delta(\rho_H+\rho_L-1)} \Delta c \right] + \gamma_{ik} \frac{\alpha\delta(\rho_H+\rho_L-1)}{1-\delta(\rho_H+\rho_L-1)} \Delta c \right\}$$

with

1.
$$\gamma_{ik} \equiv \frac{\pi_k f_H(w_{ik}) h_H(w_{ik})}{\pi_k f_H(w_{ik}) h_H(w_{ik}) + (1-\pi_k) f_L(w_{ik}) h_L(w_{ik})}$$

where $h_\theta(\cdot)$ is the pdf of the 1^{st} order statistic among $N-1$ i.i.d random variables distributed according to F_θ, $\forall \theta \in \{H, L\}$;

2. $\Delta c = c_H - c_L$ *with* $c_\theta \equiv \int_0^1 x g_\theta(x)\,\mathrm{d}x$.

The equilibrium bidding function $b^*(w_{ik}, \pi_k)$ can be decomposed in a private value (PV) and a resale value (RV) component.

$$\begin{aligned} \text{PV} &= \tfrac{w_{ik}}{1-\delta+\alpha\delta} \\ \text{RV} &= \tfrac{\alpha}{1-\delta+\alpha\delta}\left[\tfrac{\delta}{1-\delta}\left(c_L + \tfrac{1-\rho_L}{1-\delta(\rho_H+\rho_L-1)}\Delta c\right) + \gamma_{ik}\tfrac{\delta(\rho_H+\rho_L-1)}{1-\delta(\rho_H+\rho_L-1)}\Delta c\right] \end{aligned}$$

The private value component is the expected discounted use value of the good until resale takes place. An increase in the expected resale horizon ($\alpha\downarrow$) increases the private value component, and it decreases the resale value one. Bidders expect to enjoy the good for a longer time, so their use value gains importance relative to the expected future resale price. The resale value component includes a constant term, and another term which depends on belief γ_{ik}. The latter depends on the public belief π_k, and on the private use value w_{ik}. The random variable w_{ik} enters in two distinct updating. First, w_{ik} is a signal on the current state of the world because it comes from the common distribution F_θ. Second, in equilibrium a winning bidder realizes that all other $N-1$ bidders had lower private use values. This last updating is analogous to the inference carried out by a winning bidder in a static common value auction. In this respect, our model may offer a dynamic micro-foundation of a static common value auction. The future resale price is at the root of the interdependence among bidders' valuations.

3.3.3 Comparative statics

In this section we carry out a few comparative statics exercises to highlight the main determinants of different price dynamics.

The sensitiveness of $b^*(w_{ik}, \pi_k)$ with respect to γ can be measured through a simple elasticity measure:

$$\eta_\gamma^b := \frac{\gamma}{b}\frac{\partial b}{\partial \gamma} = \frac{\gamma\frac{\alpha\delta(\rho_H+\rho_L-1)}{1-\delta(\rho_H+\rho_L-1)}\Delta c}{w_{ik} + \frac{\alpha\delta}{1-\delta}\left[c_L + \frac{1-\rho_L}{1-\delta(\rho_H+\rho_L-1)}\Delta c\right] + \gamma\frac{\alpha\delta(\rho_H+\rho_L-1)}{1-\delta(\rho_H+\rho_L-1)}\Delta c}$$

A higher value of η_γ^b denotes a greater sensitivity of bidders' strategies to their present beliefs about the state of the world. It is easy to show that η_γ^b is increasing

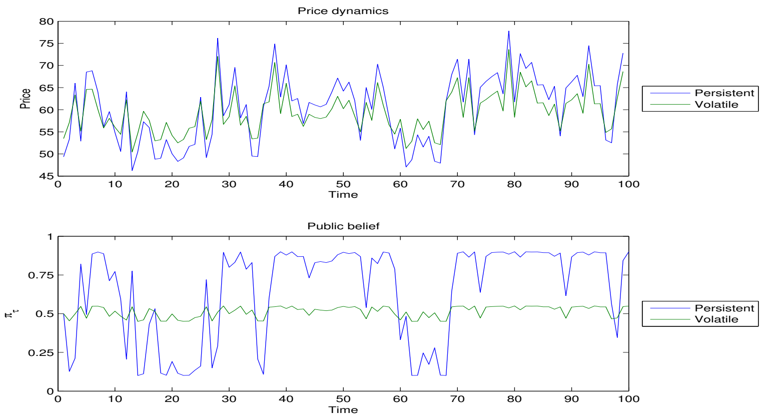

Figure 3.1: Difference in state persistence.

Simulation with $N = 4$, $\delta = 0.9$, $\alpha = 0.6$, $\pi_0 = 0.5$. The persistent line refers to $\rho_H = \rho_L = 0.9$ while the volatile one to $\rho_H = \rho_L = 0.55$. Private values w_{ik} are randomly sampled from lognormal distributions with mean $\mu_H = 8$, $\mu_L = 4$ and variance $\sigma = 2$.

in α, Δc and ρ_j, $j = H, L$. In words, the bidding strategy is more sensitive to new information if: (i) the resale horizon is shorter ($\alpha \uparrow$); (ii) the expected difference between aggregates states is larger[17] ($\Delta c \uparrow$); or (iii) each state is more persistent ($\rho_j \uparrow$). The intuition for each variable is pretty straightforward. When the resale horizon is shorter, present information is more accurate to predict the state of the world at the future time of resale. Similarly, when states of the world are more persistent, current beliefs are more precise in predicting future states. As a result, prices respond more to new information (Figure 3.1). Finally, a greater difference Δc increases the variability in the possible resale values between the two aggregate states, and agents adjust their bids more sharply.

Lastly, we derive a statistical measure of dispersion for *realized* prices. Our variance measure is derived assuming a deterministic resale horizon q, say, periods long, and a future state of the world θ_{k+1}.[18] Public beliefs move between any two trading periods according to the law of motion in Definition 3.3.1, and—for a given π_k and a fixed resale horizon q— it is immediate to get the value of π_{k+1}.[19]

[17] Specifically it is the difference in the expected second highest use value out of N bidders between the high and low state of the world.

[18] Notice the difference with the variance computed according to the subjective belief of a bidders in auction k. In this case, bidders do not know neither the present nor the future state.

[19] If we did not condition on a fixed resale horizon, we could have alternatively computed a measure of *expected* variance using as weights the probability to resale in a given future period.

Three different factors contribute to price variability:

$$\mathrm{Var}_\theta\left(b^*\left(w^{(2)}_{k+1},\pi_{k+1}\right)\Big|\pi_{k+1}\right)=\left(\tfrac{1}{1-\delta+\alpha\delta}\right)^2\Bigg[\mathrm{Var}_\theta\left(w^{(2)}\right)+\left(\tfrac{\alpha\delta(\rho_H+\rho_L-1)}{1-\delta(\rho_H+\rho_L-1)}\right)^2\mathrm{Var}_\theta\left(\gamma\left(w^{(2)}_{k+1},\pi_{k+1}\right)\Big|\pi_{k+1}\right)$$
$$+\;2\tfrac{\alpha\delta(\rho_H+\rho_L-1)}{1-\delta(\rho_H+\rho_L-1)}\,\mathrm{cov}_\theta\left(w^{(2)}_{k+1},\gamma\left(w^{(2)}_{k+1},\pi_{k+1}\right)\Big|\pi_{k+1}\right)\Bigg] \tag{3.3}$$

The first term $\mathrm{Var}_\theta\left(w^{(2)}\right)$ captures the heterogeneity in private use values. This idiosyncratic component depends on the initial distribution F_θ, and on the number of bidders N. A higher dispersion in subjective use values increases this quantity. The effect of an increase in N is not obvious, and it depends on the specific $F_\theta(\cdot)$ (see Papadatos (1995)). The second term in equation (3.3) reflects the uncertainty over the future beliefs held by the second highest bidder in auction $k+1$. It is a product of two quantities: a multiplicative constant, and the variance of γ_{k+1} conditional on π_{k+1}. The former is increasing in α and ρ_j; the latter is a complex quantity to analyze without additional assumptions on the functional forms for the pdfs. Lastly, the third term captures bidders' updating of γ_{ik} with the private use value $w_{i,k+1}$. The latter is used as an informative signal on the underlying aggregate state. The covariance term is always positive and it further increases price variability. The last two terms in equation (3.3) represent the volatility due to the uncertainty over future market conditions.

A decrease in the resale horizon ($\alpha\uparrow$) increases aggregate variability, decreasing the idiosyncratic one. The overall effect is ambiguous. A increase in the state persistence ($\rho_j\uparrow$) does not affect idiosyncratic variance, but it increases the aggregate one. Unfortunately, it is difficult to derive additional comparative statics results without assuming a specific distribution. Nonetheless, thanks to Proposition 3.3.1, it is straightforward to simulate any quantity of interest once we assume a specific F_θ.

3.4 Conclusion

This paper proposes a model for durable goods markets. We explicitly consider the possibility to re-sell an object, and we discuss what are the potential implications for learning and price dynamics.

We first present two results on the dynamics of public beliefs. First, the finer is the information publicly revealed in equilibrium, the faster is the convergence of public beliefs to the true state of the world. Second, trade protocols may lead public beliefs to move upward or downward at different rates. In particular, if only winning bids are disclosed, beliefs tend to adjust more rapidly when aggregate demand is low.

In the second part of the paper, we consider a dynamic auction model. Thanks to an analytic characterization of the bidding strategy, we provide some comparative statics results. A longer expected resale horizon increases the importance of private use values, and prices are less sensitive to current information. In this case, price volatility is mainly driven by the idiosyncratic tastes of users. If states of the world tend to last longer, prices respond more to current information. This is also the case when the difference in market conditions between high and low states is large.

This paper assumes an exogenous resale decision which is independent from previous price dynamics. This is clearly a strong assumption. Endogenous resale decisions play a decisive role in shaping market dynamics. For example, there is strong empirical evidence on the positive correlation between volume and prices in the real-estate market. Solving a dynamic auction model with endogenous entry is a challenging future research direction, and we hope to address it in the future.

3.5 Appendix

Proof Lemma 3.2.1. We prove the statement only for $\theta = H$ as an analogous arguments holds for $\theta = L$.[20] Equation (3.2) can be easily rewritten as:

$$\mathbb{E}_H^X\left[\ln \frac{f_L^{T_1} f_H^{T_2}}{f_H^{T_1} f_L^{T_2}}\right] \leq 0$$

Jensen's inequality implies:

$$\mathbb{E}_H^X\left[\ln \frac{f_L^{T_1} f_H^{T_2}}{f_H^{T_1} f_L^{T_2}}\right] \leq \ln \mathbb{E}_H^X\left[\frac{f_L^{T_1} f_H^{T_2}}{f_H^{T_1} f_L^{T_2}}\right]$$

Notice that:

$$\mathbb{E}_H^X\left[\frac{f_L^{T_1} f_H^{T_2}}{f_H^{T_1} f_L^{T_2}}\right] = \int\limits_{S_X} \frac{f_L^{T_1}(T_1(x)) f_H^{T_2}(T_2(x))}{f_H^{T_1}(T_1(x)) f_L^{T_2}(T_2(x))} f_H^X(x)\,\mathrm{d}\mu(x) = \int\limits_{S_{T_1}} \int\limits_{\mathscr{C}_{T_1}(y)} \frac{f_L^{T_1}(T_1(x)) f_H^{T_2}(T_2(x))}{f_H^{T_1}(T_1(x)) f_L^{T_2}(T_2(x))} f_H^X(x)\,\mathrm{d}\mu(x)\mathrm{d}\mu(y)$$

For every $y \in S_{T_1}$ the function $\frac{f_L^{T_1}(T_1(x)) f_H^{T_2}(T_2(x))}{f_H^{T_1}(T_1(x)) f_L^{T_2}(T_2(x))}$ is constant for every element in $\mathscr{C}_{T_1}(y)$. For $\frac{f_L^{T_1}}{f_H^{T_1}}$ this is true by definition of $\mathscr{C}_{T_1}$, while it follows from coarseness for $\frac{f_H^{T_2}}{f_L^{T_2}}$. Then,

$$\int\limits_{S_{T_1}} \int\limits_{\mathscr{C}_{T_1}(y)} \frac{f_L^{T_1}(T_1(x)) f_H^{T_2}(T_2(x))}{f_H^{T_1}(T_1(x)) f_L^{T_2}(T_2(x))} f_H^X(x)\,\mathrm{d}\mu(x)\mathrm{d}\mu(y) = \int\limits_{S_{T_1}} \frac{f_L^{T_1}(y) f_H^{T_2}(y)}{f_H^{T_1}(y) f_L^{T_2}(y)} f_H^{T_1}(y)\,\mathrm{d}\mu(y) = \int\limits_{S_{T_1}} \frac{f_H^{T_2}(y)}{f_L^{T_2}(y)} f_L^{T_1}(y)\,\mathrm{d}\mu(y)$$

[20] This proof uses coarseness in order to reduce the expression to a standard Gibbs' inequality.

As $\sigma(T_2) \subset \sigma(T_1)$ we can rewrite $f_\theta^{T_2}(z) = \int\limits_{y \in S_{T_1}: T_2(y)=z} f_\theta^{T_1}(y) \mathrm{d}\mu(y)$. Therefore:

$$\int\limits_{S_{T_1}} \frac{f_H^{T_2}(y)}{f_L^{T_2}(y)} f_L^{T_1}(y)\, \mathrm{d}\mu(y) = \int\limits_{S_{T_2}} \frac{f_H^{T_2}(z)}{f_L^{T_2}(z)} \int\limits_{y \in S_{T_1}: T_2(y)=z} f_L^{T_1}(y) \mathrm{d}\mu(y) \mathrm{d}\mu(z) = \int\limits_{S_{T_2}} \frac{f_H^{T_2}(z)}{f_L^{T_2}(z)} f_L^{T_2}(z) \mathrm{d}\mu(z) = 1$$

As a result $\ln \mathbb{E}_H^X \left[\frac{f_L^{T_1} f_H^{T_2}}{f_H^{T_1} f_L^{T_2}} \right] = \ln 1 = 0$.

Lastly, observe that Jensen's inequality holds strictly. In fact, ln is a strictly concave function, and $\frac{f_L^{T_1} f_H^{T_2}}{f_H^{T_1} f_L^{T_2}}$ is not constant almost everywhere because coarseness implies the existence at least two sets A, B s.t. $A \subset B$, $A \notin \sigma(T_2)$ and $\mu(A) > 0$ where $T_2(x)$ is constant $\forall x \in B$ while $T_1(x) \neq T_1(x')$ for $x \in A$ and $x' \in B \backslash A$. ■

Proof Lemma 3.2.2. As $\pi_0 = \frac{1}{2}$ we have $l_k^T = \sum_{i=0}^{k-1} \Delta l_{i+1}^T$ where $\Delta l_{i+1}^T = l_{i+1}^T - l_i^T = \ln \frac{f_H^T(y_i)}{f_L^T(y_i)}$ $i = 0, 1, ..., k$ is a sequence of i.i.d. random variables.

Hitting times τ_H and τ_L can be equivalently stated in terms l_k:

$$\tau_H^l := \inf \left\{ k > 0 : l_k^T \geq \ln \frac{1-\varepsilon}{\varepsilon} \right\} \qquad \tau_L^l := \inf \left\{ k > 0 : l_k^T \leq \ln \frac{\varepsilon}{1-\varepsilon} \right\}$$

Applying Wald (1944) lemma to the sequence of i.i.d random variables Δl_i:

$$\mathbb{E}_\theta[l_{\tau_\theta^l}^T] = \mathbb{E}_\theta[\tau_\theta^l] \mathbb{E}_\theta[\Delta l^T] \qquad \forall \theta \in \{H, L\} \tag{3.4}$$

By Gibbs' inequality $\mathbb{E}_H[\Delta l^T] > 0$ and $\mathbb{E}_L[\Delta l^T] < 0$. If $\mathbb{E}_H[\Delta l^T] + \mathbb{E}_L[\Delta l^T] < 0$ then:

$$\mathbb{E}_L[\Delta l^T] = -\left(\mathbb{E}_H[\Delta l^T] + c \right)$$

where $c \equiv -\int_{S_T} \ln \frac{f_H^T(y)}{f_L^T(y)} \left(f_H^T(y) + f_L^T(y) \right) \mathrm{d}\mu(y) > 0$. Note that c only depends on the primitives and it is independent of ε.

Substituting in equation (3.4):

$$\begin{aligned} \mathbb{E}_H[\tau_H^l] \mathbb{E}_H[\Delta l^T] = {} & \mathbb{E}_H[l_{\tau_H^l}^T] \\ \mathbb{E}_L[\tau_L^l](\mathbb{E}_H[\Delta l^T] + c) = {} & -\mathbb{E}_L[l_{\tau_L^l}^T] \end{aligned}$$

Hence:

$$\begin{aligned} \mathbb{E}_H[\Delta l^T](\mathbb{E}_H[\tau_H^l] - \mathbb{E}_L[\tau_L^l]) = {} & \mathbb{E}_H[l_{\tau_H^l}^T] + \mathbb{E}_L[l_{\tau_L^l}^T] + \mathbb{E}_L[\tau_L^l] c \\ \Longrightarrow \qquad \mathbb{E}_H[\Delta l^T] \left(\tfrac{\mathbb{E}_H[\tau_H^l]}{\mathbb{E}_L[\tau_L^l]} - 1 \right) = {} & \frac{\mathbb{E}_H[l_{\tau_H^l}^T] + \mathbb{E}_L[l_{\tau_L^l}^T]}{\mathbb{E}_L[\tau_L^l]} + c \end{aligned}$$

Note that $|\Delta l^T| < M$ implies $\mathbb{E}_L[l_{\tau_L^l}^T] \geq \ln \frac{\varepsilon}{1-\varepsilon} - M$ and by definition $\mathbb{E}_H[l_{\tau_H^l}^T] \geq \ln \frac{1-\varepsilon}{\varepsilon}$, hence:

$$\mathbb{E}_H[\Delta l^T] \left(\frac{\mathbb{E}_H[\tau_H^l]}{\mathbb{E}_L[\tau_L^l]} - 1 \right) \geq \frac{-M}{\mathbb{E}_L[\tau_L^l]} + c$$

On the other hand, by Wald lemma:

$$\mathbb{E}_L[\tau_L^l] = \frac{\mathbb{E}_L[l^T_{\tau_L^l}]}{\mathbb{E}_L[\Delta l^T]} \geq \frac{\ln\frac{1-\varepsilon}{\varepsilon}}{M} > 0$$

so

$$\mathbb{E}_H[\Delta l^T](\frac{\mathbb{E}_H[\tau_H^l]}{\mathbb{E}_L[\tau_L^l]} - 1) \geq -\frac{M^2}{\ln\frac{1-\varepsilon}{\varepsilon}} + c$$

Since $0 < \mathbb{E}_H[\Delta l^T] < M$:

$$\begin{aligned}
\frac{\mathbb{E}_H[\tau_H^l]}{\mathbb{E}_L[\tau_L^l]} - 1 > & \quad -\frac{M}{\ln\frac{1-\varepsilon}{\varepsilon}} + \frac{c}{\mathbb{E}_H[\Delta l^T]} \\
= & \quad -\frac{M}{\ln\frac{1-\varepsilon}{\varepsilon}} + \frac{-\mathbb{E}_H[\Delta l^T] - \mathbb{E}_L[\Delta l^T]}{\mathbb{E}_H[\Delta l^T]} \\
= & \quad -\frac{M}{\ln\frac{1-\varepsilon}{\varepsilon}} - 1 + \left|\frac{\mathbb{E}_L[\Delta l^T]}{\mathbb{E}_H[\Delta l^T]}\right| \text{ as } \mathbb{E}_L[\Delta l^T] < 0
\end{aligned}$$

Note that $\ln\frac{1-\varepsilon}{\varepsilon} \to \infty$ as $\varepsilon \to 0$, hence $\forall \delta > 0$, $\exists \overline{\varepsilon} > 0$ such that $\forall \varepsilon < \overline{\varepsilon}$, $\frac{M}{\ln\frac{1-\varepsilon}{\varepsilon}} < \delta$. The proof for the other case is symmetric. ■

Proof Proposition 3.3.1.

Consider the subgame starting from auction k. Notice that Δ_k is statistically independent of the underlying state and all private values, hence we can integrate it out when calculating the expected payoff of bidder ik:

$$\begin{aligned}
& \mathbb{E}[u_{ik}(b, b^*_{-ik}; w_{ik})|w_{ik}, \pi_k] \\
= \quad & \mathbb{P}(b_k^{*(2)} < b|w_{ik}, \pi_k)\Big\{-\mathbb{E}(b_k^{*(2)}|w_{ik}, \pi_k) \\
+ \quad & \sum_{x=1}^{\infty} \alpha(1-\alpha)^{x-1}\left[\sum_{s=0}^{x-1} \delta^s w_{ik} + \delta^x \mathbb{E}(b_{k+1}^{*(2)}|w_{ik}, \pi_k, \Delta_k = x)\right]\Big\} \\
= \quad & \mathbb{P}(b_k^{*(2)} < b|w_{ik}, \pi_k)\left\{-\mathbb{E}(b_k^{*(2)}|w_{ik}, \pi_k) + \sum_{x=1}^{\infty} \alpha(1-\alpha)^{x-1}\frac{1-\delta^x}{1-\delta} w_{ik}\right. \\
+ \quad & \left.\sum_{x=1}^{\infty}\left[\frac{\alpha}{1-\alpha}(\delta - \alpha\delta)^x\, \mathbb{E}(b_{k+1}^{*(2)}|w_{ik}, \pi_k, \Delta_k = x)\right]\right\} \\
= \quad & \mathbb{P}(\theta_{t_k} = H|w_{ik}, \pi_k)\mathbb{P}(b_k^{*(2)} < b|\theta_{t_k} = H)\Big\{-\mathbb{E}(b_k^{*(2)}|\theta_{t_k} = H) + \sigma w_{ik} \\
+ \quad & \sum_{x=1}^{\infty}\left[\frac{\alpha}{1-\alpha}(\delta - \alpha\delta)^x E(b_{k+1}^{*(2)}|\theta_{t_k} = H, \Delta_k = x)\right]\Big\} \\
+ \quad & \mathbb{P}(\theta_{t_k} = L|w_{ik}, \pi_k)\mathbb{P}(b_k^{*(2)} < b|\theta_{t_k} = L)\Big\{-\mathbb{E}(b_k^{*(2)}|\theta_{t_k} = L) + \sigma w_{ik} \\
+ \quad & \sum_{x=1}^{\infty}\left[\frac{\alpha}{1-\alpha}(\delta - \alpha\delta)^x\, \mathbb{E}(b_{k+1}^{*(2)}|\theta_{t_k} = L, \Delta_k = x)\right]\Big\}
\end{aligned}$$

For convenience, let us introduce the following notation:

$z_k \equiv w_k^{(2)} \equiv v_{t_k}^{(2)}$:	2^{nd} highest realization of private values in auction k
$\rho_j^x \equiv ((P^x)_{j1}, (P^x)_{j2})$	the j-th row of matrix P^x, $j \in \{1,2\}, \forall x \in \mathbb{N}^+$
$e_{k+1} \equiv (e_{k+1}^H, e_{k+1}^L)^\intercal$ $e_{k+1}^\theta \equiv \mathbb{E}(b_{k+1}^{*(2)} \vert \theta_{t_{k+1}} = \theta)$	the expectation of equilibrium resale revenue conditional on the state of next auction
$\widetilde{\gamma}_k = (\widetilde{\gamma}_k, 1-\widetilde{\gamma}_k)$ $\widetilde{\gamma}_k \equiv \frac{\pi_k f_H(z_k) h_H(z_k)}{\pi_k f_H(z_k) h_H(z_k) + (1-\pi_t) f_L(z_k) h_L(z_k)}$	the belief used by the 2nd-highest-value bidder of auction k in his bidding function.

Using the notations above, we have:

$$\mathbb{E}(b_{k+1}^{*(2)}|\theta_{t_k} = H, \Delta_k = x) = \rho_1^x e_{k+1};$$
$$\mathbb{E}(b_{k+1}^{*(2)}|\theta_{t_k} = L, \Delta_k = x) = \rho_2^x e_{k+1}.$$

Assuming monotone and symmetric bidding strategy, we can rewrite bidder ik's problem as:

$$\begin{aligned}\max_b \quad & \pi_{ik} \int_0^{b^{*-1}(b;\pi_k)} \left[-b^*(y;\pi_k) + \sigma w_{ik} + \sum_{x=1}^{\infty} \frac{\alpha}{1-\alpha}(\delta - \alpha\delta)^x \rho_1^x e_{k+1}\right] h_H(y)\,\mathrm{d}y \\ + \quad & (1-\pi_{ik}) \int_0^{b^{*-1}(b;\pi_k)} \left[-b^*(y;\pi_k) + \sigma w_{ik} + \sum_{x=1}^{\infty} \frac{\alpha}{1-\alpha}(\delta - \alpha\delta)^x \rho_2^x e_{k+1}\right] h_L(y)\,\mathrm{d}y\end{aligned}$$

where $\pi_{ik} \equiv \mathbb{P}(\theta_{t_k} = H|w_{ik}, \pi_k) = \frac{\pi_k f_H(w_{ik})}{\pi_k f_H(w_{ik}) + (1-\pi_k) f_L(w_{ik})}$, bidder ik's posterior about θ_{t_k}. Note that e_{k+1}, the expected equilibrium resale revenue, will depend on x, the realization of Δ_k, and y, the realization of z_k, through public belief π_{k+1}, therefore it cannot be taken out of the integral.

FOC yields:

$$\begin{aligned}0 \quad = \quad & \pi_{ik} \left[-b^*(w_{ik};\pi_k) + \sigma w_{ik} + \sum_{x=1}^{\infty} \frac{\alpha}{1-\alpha}(\delta - \alpha\delta)^x \rho_1^x e_{k+1}\right] h_H(w_{ik}) \\ + \quad & (1-\pi_{ik}) \left[-b^*(w_{ik};\pi_k) + \sigma w_{ik} + \sum_{x=1}^{\infty} \frac{\alpha}{1-\alpha}(\delta - \alpha\delta)^x \rho_2^x e_{k+1}\right] h_L(w_{ik})\end{aligned}$$

Using γ_{it} and B defined in the proposition we can rewrite the FOC as

$$\begin{aligned}b^*(w_{ik};\pi_k) \quad & = \quad \sigma w_{ik} + \gamma_{ik} \left[\sum_{x=1}^{\infty} \frac{\alpha}{1-\alpha}(\delta - \alpha\delta)^x P^x e_{k+1}\right] \\ & = \quad \sigma w_{ik} + \left[\sum_{x=1}^{\infty} \frac{\alpha}{1-\alpha}(\delta - \alpha\delta)^x (\gamma_{ik} P^x) e_{k+1}\right]\end{aligned}$$

Now we need to solve for the equilibrium object e_{k+1}.

Since the bidding strategy is time-invariant:

$$\begin{aligned} e_{k+1}^{\theta} &\equiv \mathbb{E}(b_{k+1}^{*(2)} \big| \theta_{t_{k+1}} = \theta) \\ &= \mathbb{E}\left[b^*(z_{k+1}; \pi_{k+1}(\pi_k, y = w_{ik}, x; P)) \big| \theta_{t_{k+1}} = \theta \right] \\ &= \mathbb{E}\left[\sigma z_{k+1} + \sum_{x'=1}^{\infty} \frac{\alpha}{1-\alpha} (\delta - \alpha\delta)^{x'} (\widetilde{\gamma}_{k+1} P^{x'}) e_{k+2} \middle| \theta_{t_{k+1}} = \theta \right] \end{aligned}$$

Put it back into the bidding function above:

$$\begin{aligned} & b^*(w_{ik}; \pi_k) \\ = \quad & \sigma w_{ik} + \left[\sum_{x=1}^{\infty} \frac{\alpha}{1-\alpha} (\delta - \alpha\delta)^x (\gamma_{ik} P^x) e_{k+1} \right] \\ = \quad & \sigma w_{ik} + \sum_{x=1}^{\infty} \frac{\alpha}{1-\alpha} (\delta - \alpha\delta)^x (\gamma_{ik} P^x) \sigma \begin{pmatrix} E[z_{k+1} | \theta_{t_{k+1}} = H] \\ E[z_{k+1} | \theta_{t_{k+1}} = L] \end{pmatrix} \\ + \quad & \sum_{x=1}^{\infty} \frac{\alpha}{1-\alpha} (\delta - \alpha\delta)^x (\gamma_{ik} P^x) \begin{pmatrix} \mathbb{E}\left[\sum_{x'=1}^{\infty} \frac{\alpha}{1-\alpha} (\delta - \alpha\delta)^{x'} (\widetilde{\gamma}_{k+1} P^{x'}) e_{k+2} \middle| \theta_{t_{k+1}} = H \right] \\ \mathbb{E}\left[\sum_{x'=1}^{\infty} \frac{\alpha}{1-\alpha} (\delta - \alpha\delta)^{x'} (\widetilde{\gamma}_{k+1} P^{x'}) e_{k+2} \middle| \theta_{t_{k+1}} = L \right] \end{pmatrix} \\ = \quad & \sigma w_{ik} + \sigma \sum_{x=1}^{\infty} \frac{\alpha}{1-\alpha} (\delta - \alpha\delta)^x \mathbb{E}_{ik}[z_{k+1}] \\ + \quad & \sum_{x=1}^{\infty} \frac{\alpha}{1-\alpha} (\delta - \alpha\delta)^x \mathbb{E}_{ik} \left[\sum_{x'=1}^{\infty} \frac{\alpha}{1-\alpha} (\delta - \alpha\delta)^{x'} (\widetilde{\gamma}_{k+1} P^{x'}) e_{k+2} \right] \end{aligned}$$

where $\mathbb{E}_{ik}[\cdot]$ denote the expectation of bidder ik conditional on the event that he wins auction k and the highest value among others is exactly equal to his value, and his waiting time for resale is x.

Now consider e_{k+2}. Let us label the 2nd highest bidder in auction k as bidder $\widetilde{k}$, $\forall k \in \mathbb{N}$.

$$\begin{aligned} e_{k+2}^{\theta} &\equiv \mathbb{E}\left[b_{k+2}^{*(2)} | \theta_{t_{k+2}} = \theta \right] \\ &= \mathbb{E}\left[b^*(z_{k+2}; \pi_{k+2}(\pi_{k+1}, y = z_{k+1}, x; P)) | \theta_{t_{k+2}} = \theta \right] \\ &= \mathbb{E}\left\{ \sigma z_{k+2} + \sum_{x''=1}^{\infty} \frac{\alpha}{1-\alpha} (\delta - \alpha\delta)^{x''} (\widetilde{\gamma}_{k+2} P^{x''}) e_{+3} \middle| \theta_{t_{k+2}} = \theta \right\} \end{aligned}$$

Hence:

$$\left(\widetilde{\gamma}_{k+1} P^{x'} \right) e_{k+2} = \sigma \mathbb{E}_{\widetilde{k+1}}[z_{k+2}] + \mathbb{E}_{\widetilde{k+1}} \left[\sum_{x''=1}^{\infty} \frac{\alpha}{1-\alpha} (\delta - \alpha\delta)^{x''} \left(\widetilde{\gamma}_{k+2} P^{x''} \right) e_{k+3} \right]$$

where $\mathbb{E}_{\widetilde{k+1}}[\cdot]$ denote the expectation of bidder $\widetilde{k+1}$ conditional on the event that he wins auction $k+1$ and the highest value among others is exactly equal to his value, and his waiting time for resale is x'.

Plugging back this value into the bidding function of bidder ik:

$$
\begin{aligned}
& b^*(w_{ik};\pi_k) \\
= \; & \sigma w_{ik} + \sigma \sum_{x=1}^{\infty} \frac{\alpha}{1-\alpha}(\delta-\alpha\delta)^x \mathbb{E}_{ik}[z_{k+1}] \\
+ \; & \sum_{x=1}^{\infty} \frac{\alpha}{1-\alpha}(\delta-\alpha\delta)^x \mathbb{E}_{ik}\left[\sum_{x'=1}^{\infty} \frac{\alpha}{1-\alpha}(\delta-\alpha\delta)^{x'} \sigma \mathbb{E}_{\widetilde{k+1}}[z_{k+2}]\right] + r_{ik} \\
= \; & \sigma w_{ik} + \sigma \sum_{x=1}^{\infty} \frac{\alpha}{1-\alpha}(\delta-\alpha\delta)^x \mathbb{E}_{ik}[z_{k+1}] \\
+ \; & \sigma \sum_{x=1}^{\infty} \frac{\alpha}{1-\alpha}(\delta-\alpha\delta)^x \left[\sum_{x'=1}^{\infty} \frac{\alpha}{1-\alpha}(\delta-\alpha\delta)^{x'} E_{ik}[z_{k+2}]\right] + r_{ik} \\
= \; & \sigma w_{ik} + \sigma \sum_{x=1}^{\infty} \frac{\alpha}{1-\alpha}(\delta-\alpha\delta)^x \gamma_{ik} P^x c \\
+ \; & \sigma \sum_{x=1}^{\infty} \frac{\alpha}{1-\alpha}(\delta-\alpha\delta)^x \left[\sum_{x'=1}^{\infty} \frac{\alpha}{1-\alpha}(\delta-\alpha\delta)^{x'} \gamma_{ik} P^{x+x'} c\right] + r_{ik} \\
= \; & \sigma w_{ik} + \sigma\gamma_{ik} \left[\sum_{x=1}^{\infty} \frac{\alpha}{1-\alpha}(\delta-\alpha\delta)^x P^x\right] c + \sigma\gamma_{ik} \left[\sum_{x=1}^{\infty} \frac{\alpha}{1-\alpha}(\delta-\alpha\delta)^x P^x\right]^2 c + r_{ik} \\
= \; & \sigma w_{ik} + \sigma\gamma_{ik} B c + \sigma\gamma_{ik} B^2 c + r_{ik}
\end{aligned}
$$

The second equation comes from law of iterated expectation and the fourth equation comes from the fact that waiting time is i.i.d. across auctions.

The residual term r_{ik} is:

$$
\sum_{x=1}^{\infty} \frac{\alpha}{1-\alpha}(\delta-\alpha\delta)^x \mathbb{E}_{ik}\left[\sum_{x'=1}^{\infty} \frac{\alpha}{1-\alpha}(\delta-\alpha\delta)^{x'} \sigma \mathbb{E}_{\widetilde{k+1}}\left[\sum_{x''=1}^{\infty} \frac{\alpha}{1-\alpha}(\delta-\alpha\delta)^{x''} (\widetilde{\gamma}_{k+2} P^{x''}) e_{k+3}\right]\right]
$$

Note that the expected present value of the resale revenue from auction $k+m$ for bidder ik goes to 0 as m goes to infinity, due to the existence of discount rate δ. Hence we can recursively solve for the bidding function following the argument above, and finally get:

$$
\begin{aligned}
b^*(w_{ik};\pi_k) &= \sigma w_{ik} + \sigma\gamma_{ik} \left(\sum_{s=1}^{\infty} B^s\right) c \\
&= \sigma[w_{ik} + \gamma_{ik} B (I-B)^{-1} c]
\end{aligned}
$$

To complete the proof, we need to verify that this is indeed a monotone bidding function.

Let $D = \begin{bmatrix} d_{11} & d_{12} \\ d_{21} & d_{22} \end{bmatrix} \equiv B(I-B)^{-1}$ and rewrite $b^*(w_{ik};\pi_k)$ as

$$\begin{aligned} b^*(w_{ik};\pi_k) &= \sigma(w_{ik} + \gamma_{ik} Dc) \\ &= \sigma\{w_{ik} + d_{21}c_H + d_{22}c_L + \gamma_{ik}[(d_{11} - d_{22} + d_{12} - d_{21})c_L + (d_{11} - d_{21})\Delta c]\} \end{aligned}$$

where $\Delta c \equiv c_H - c_L > 0$ and d_{ij}, $i,j = 1,2$, to be determined.

The matrix $B = \alpha\delta P[I - \delta(1-\alpha)P]^{-1}$ is equal to:

$$\begin{aligned} B &= \alpha\delta \begin{bmatrix} \rho_H & 1-\rho_H \\ 1-\rho_L & \rho_L \end{bmatrix} \begin{bmatrix} 1-\delta(1-\alpha)\rho_H & \delta(1-\alpha)(1-\rho_H) \\ -\delta(1-\alpha)(1-\rho_L) & 1-\delta(1-\alpha)\rho_L \end{bmatrix}^{-1} \\ &= \frac{\alpha\delta}{(1-\delta(1-\alpha))(1-\delta(1-\alpha)(\rho_H+\rho_L-1))} \cdot \\ &\quad \begin{bmatrix} \rho_H - \delta(1-\alpha)(\rho_H+\rho_L-1) & 1-\rho_H \\ 1-\rho_L & \rho_L - \delta(1-\alpha)(\rho_H+\rho_L-1) \end{bmatrix} \equiv \kappa \begin{bmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{bmatrix} \end{aligned}$$

Therefore:

$$\begin{aligned} D &= B(I-B)^{-1} = \kappa \begin{bmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{bmatrix} \begin{bmatrix} 1-\kappa b_{11} & -\kappa b_{12} \\ -\kappa b_{21} & 1-\kappa b_{22} \end{bmatrix}^{-1} \\ &= \frac{1}{(1-\kappa b_{11})(1-\kappa b_{22}) - \kappa^2 b_{12}b_{21}} \cdot \begin{bmatrix} \kappa b_{11}(1-\kappa b_{22}) + \kappa^2 b_{12}b_{21} & \kappa b_{12} \\ \kappa b_{21} & \kappa b_{22}(1-\kappa b_{11}) + \kappa^2 b_{12}b_{21} \end{bmatrix} \end{aligned}$$

Plugging back the value of κ and $b_{ij}, i,j = 1,2$ we have:

$$\begin{aligned} (1-\kappa b_{11})(1-\kappa b_{22}) - \kappa^2 b_{12}b_{21} &= \tfrac{(1-\delta)(1-\delta(\rho_H+\rho_L-1))}{(1-\delta+\alpha\delta)(1-\delta(1-\alpha)(\rho_H+\rho_L-1))}; \\ \kappa b_{11}(1-\kappa b_{22}) + \kappa^2 b_{12}b_{21} &= \tfrac{\alpha\delta}{(1-\delta+\alpha\delta))(1-\delta(1-\alpha)(\rho_H+\rho_L-1))}(\rho_H - \delta(\rho_H+\rho_L-1)); \\ \kappa b_{22}(1-\kappa b_{11}) + \kappa^2 b_{12}b_{21} &= \tfrac{\alpha\delta}{(1-\delta+\alpha\delta)(1-\delta(1-\alpha)(\rho_H+\rho_L-1))}(\rho_L - \delta(\rho_H+\rho_L-1)); \\ \kappa b_{12} &= \tfrac{\alpha\delta}{(1-\delta+\alpha\delta)(1-\delta(1-\alpha)(\rho_H+\rho_L-1))}(1-\rho_H); \\ \kappa b_{21} &= \tfrac{\alpha\delta}{(1-\delta+\alpha\delta)(1-\delta(1-\alpha)(\rho_H+\rho_L-1))}(1-\rho_L) \end{aligned}$$

Therefore:

$$D = \frac{\alpha\delta}{(1-\delta)(1-\delta(\rho_H+\rho_L-1))} \begin{bmatrix} \rho_H - \delta(\rho_H+\rho_L-1) & 1-\rho_H \\ 1-\rho_L & \rho_L - \delta(\rho_H+\rho_L-1) \end{bmatrix}$$

If we plug back the elements of D into the bidding function we get:

$$b^*(w_{ik};\pi_k) = \sigma\left\{w_{ik} + \frac{\alpha\delta}{1-\delta}\left[c_L + \frac{1-\rho_L}{1-\delta(\rho_H+\rho_L-1)}\Delta c\right] + \gamma_{ik}\frac{\alpha\delta(\rho_H+\rho_L-1)}{1-\delta(\rho_H+\rho_L-1)}\Delta c\right\}$$

Since $\rho_H + \rho_L \in [1,2]$ and γ_{ik} is strictly monotone in w_{ik}, $b^*(w_{ik};\pi_k)$ is clearly strictly monotone in w_{ik} as well.

Lastly, notice that the $b^*(w_{ik};\pi_k)$ is indeed the unique solution to the FOC, which implies that the expected payoff of each bidder would be a single-peaked function of her bid, hence FOC is sufficient for optimality. ■

Bibliography

HAILE, P.A. (2000): "Partial Pooling at the Reserve Price in Auctions with Resale Opportunities," *Games and Economic Behaviour*, 33, 231–248.

HAILE, P.A (2001): "Auctions with Resale Markets: An Application to U.S. Forest Service Timber Sales," *American Economic Review*, 91, 399–427.

HAILE, P. A. (2003): "Auctions with private uncertainty and resale opportunities," *Journal of Economic Theory*, 108, 72–110.

HALAFIR, I. E. AND V. KRISHNA (2008): "Asymmetric Auctions with Resale," *American Economic Review*, 98, 87–112.

HAYEK, F. A. (1945): "The Use of Knowledge in Society," *American Economic Review*, 35, 519–530.

GARRATT, R. AND T. TROĞER (2006): "Speculation in Standard Auctions with Resale," *Econometrica*, 74, 753–770.

GARRATT, R., T. TROĞER, AND C. ZHENG (2009): "Collusion via Resale," *Econometrica*, 77, 1095–1136.

GUPTA, M. AND B. LEBRUN (1999): "First price auctions with resale," *Economics Letters*, 64, 181–185.

JU, N., AND J. MIAO (2012): "Ambiguity, Learning and Asset Returns," *Econometrica*, 80, 559–591.

LEBRUN, B. (2012): "Optimality and the English and second-price auctions with resale," *Games and Economic Behaviour*, 75, 731–751.

MILGROM, P. R., AND R. J. WEBER (1982): "A Theory of Auctions and Competitive Bidding," *Econometrica*, 50, 1089–1122.

MYERSON, R. B. (1981): "Optimal Auction Design," *Mathematics of Operations Research*, 6, 58–73.

PAPADATOS, N. (1995): “Maximum Variance of Order Statistics,” *Ann. Inst. Statist. Math*, 47, 185–193.

TIMMERMANN, A. (1993): “How Learning in Financial Markets Generates Excess Volatility and Predictability in Stock Prices,” *Quarterly Journal of Economics*, 108, 1135–1145.

VIRAG, G. (2013): “First-price auctions with resale: the case of many bidders,” *Economic Theory*, 52, 129–163.

WALD, A. (1944): “On cumulative sums of random variables, ” *The Annals of Mathematical Statistics*, 15, 283–296.

WEITZMAN, M. L. (2007): “Subjective Expectations and Asset-Return Puzzles,” *American Economic Review*, 97, 1102–1130.

ZHENG, C. Z. (2002): “Optimal Auction with Resale,” *Econometrica*, 70, 2197–2224.