Cycling and Categorical Learning in Decentralized Adverse Selection Economies

Philippe Jehiel, Erik Mohlin

WORKING PAPER N° 2022 – 24

Cycling and Categorical Learning in Decentralized Adverse Selection Economies

Philippe Jehiel
Erik Mohlin

JEL Codes: C70, C73, D82, D83, D91
Keywords: Adverse selection; Bounded rationality; Categorization; Learning; Model misspeciÖcation; OTC markets

Cycling and Categorical Learning in Decentralized Adverse Selection Economies *

Philippe Jehiel[†] Erik Mohlin[‡]

September 2, 2021

Abstract

We study learning in a decentralized pairwise adverse selection economy, where buyers have access to the quality of traded goods but not to the quality of non-traded goods. Buyers categorize ask prices in order to predict quality as a function of ask price. The categorization is endogenously determined so that outcomes that are observed more often are categorized more finely, and within each category beliefs reflect the empirical average. This leads buyers to have a very fine understanding of the relationship between qualities and ask prices for prices below the current market price, but only a coarse understanding above that price. We find that this induces a price cycle involving the Nash equilibrium price, and one or more higher prices.

Keywords: Adverse selection; Bounded rationality; Categorization; Learning; Model misspecification; OTC markets.

JEL codes: C70, C73, D82, D83, D91.

1 Introduction

Consider an adverse selection market of the Akerlof type in which the sellers know the quality of their good, but the buyers do not. It is well known that if buyers are fully rational such adverse selection may lead to low volumes of trade (Akerlof 1970). However,

*This paper has benefited from comments by Drew Fudenberg, Topi Miettinen, Alexandros Rigos, and Larry Samuelson. We also thank participants at LEG19 at Bar-Ilan University for comments. Alexandros Rigos provided very valuable assistance with computations. Erik Mohlin is grateful for financial support from the Swedish Research Council (grant 2015-01751) and the Knut and Alice Wallenberg Foundation (Wallenberg Academy Fellowship 2016-0156). Philippe Jehiel thanks the European Research Council (grant no. 742816) for funding.

†Paris School of Economics and University College London. Address: PSE, 48 boulevard Jourdan, 75014 Paris, France. E-mail: jehiel@enpc.fr.

‡Lund University and the Institute for Futures Studies (Stockholm). Address: Lund University Department of Economics, Tycho Brahes väg 1, 220 07 Lund, Sweden. E-mail: erik.mohlin@nek.lu.se.

the kind of inference required for such a prediction has sometimes been challenged in light of the counterfactual reasoning it may involve. In particular, if one only observes the quality of past traded goods, and not the quality of non-traded goods, it may be hard to infer what the quality would have been if the transaction prices had been higher than those observed from previous trades.

We consider such adverse selection markets where pairs consisting of a buyer and a seller are matched to have the opportunity to trade with each other. The trading mechanism takes the form of a double auction in which, in each pair, the seller submits an ask price and the buyer submits a bid price and a transaction takes place only if the bid price exceeds the ask price (at a price assumed to be equal to the bid price to make the analysis simpler on the seller's side). Such a trading procedure can be viewed as decentralized in the sense that the trading conditions are decided for each pair separately. We note that over-the-counter (OTC) markets are typically decentralized (a limited number of sellers and buyers, out of the total population of buyers and sellers, meet to negotiate and transact). OTC markets are important both in terms of volume of trade and in recent debates about policy and regulation (Weill 2020). Our analysis will be used to shed a novel light on fluctuations in OTC markets. By comparison, we will also discuss how the analysis changes if centralized trading mechanisms are considered instead.

We envision the following dynamic setting. Time is discrete and in every period a new generation of buyers and sellers are matched to interact with each other. The buyers at time t have information about all bid and ask prices from period $t-1$ as well as information about quality in matches where trade occurred. Crucially, we assume that the quality of non-traded goods is not observed. Based on the data observed from the previous period, buyers (who are uninformed of the quality) form beliefs about how the behavior of sellers (the ask prices) relate to the quality of the proposed good.[1] The focus on recent observations may be due to limited memory or limited availability of historical data, and is in line with the well-documented recency bias observed in experiments on learning (Agarwal et al.. 2008; Erev and Haruvy 2014).[2]

We assume that belief formation is based on categorical thinking. Buyers bundle ask prices into categories and believe that the quality associated with an ask price in a particular category is equal to the average quality of objects with ask prices in the same category, as observed in the previous period (similar to the modeling of analogy based expectation equilibrium (ABEE) of Jehiel 2005 and Jehiel and Koessler 2008). Importantly, we also assume that the number and size of categories depend endogenously on the available data in agreement with the bias–variance trade-off routinely considered

[1] We discuss below the effect of longer memories.

[2] From an evolutionary perspective recency bias may reflect an adaptation to changing environments where old data are generally less informative.

in statistics. This implies that each category should contain approximately the same mass of transactions (see Mohlin 2014 for an explicit model).

This heuristic method to form beliefs has important implications for the analysis of adverse selection markets. We assume that a small minority of informed buyers are able to observe the quality of the seller they are matched with, while the remaining uninformed buyers do not observe the quality of the good of their seller and form their beliefs according to the heuristic method. Suppose that in pairs involving uninformed buyers, the bid price was p^* in the previous period. Then in those matches, only qualities for which the ask price is lower than p^* will be observed. Thus, it will not be possible from those matches (which are more numerous than matches with ask price above p^*) to relate ask prices above p^* to quality. Instead, such a relationship will have to be derived from the matches involving informed buyers.[3] Since informed buyers are few, the categorization will be much coarser for ask prices above p^* than for ask prices below p^* (for which data are abundant thanks to the pairs involving uninformed buyers). In turn, the coarse belief for ask prices above p^* will lead buyers to have an excessively optimistic view of the effect of increasing slightly the bid price above p^*. We pay special attention to the limiting case where the number of categories below p^* goes to infinity, while there are only a few categories above p^*. This gives the buyers a perfect understanding of how quality is related to the ask price for prices below p^*, as in the standard Bayes–Nash analysis, but only a coarse understanding of how quality is related to the ask price for prices above p^*.

Such an erroneous perception will cause the prices to follow a cyclical pattern. Starting from a price that corresponds to the Bayes–Nash equilibrium, we will establish that the bid price chosen by uninformed buyers will have to be strictly larger in the next generation. The price chosen by uninformed buyers may then increase over the next few generations, but at some point when it gets too high, increasing the bid price will not look profitable and at this point uninformed buyers will quote the Nash equilibrium bid price instead. The bid prices of uninformed buyers will cycle from then on, always being weakly above the price arising in the full rationality benchmark.

The key to understanding why our heuristic leads to cycles is to realize that if the bid price of uninformed buyers becomes very large, then the next generation will have all the needed data to compute the correct best response, which is the Nash equilibrium price. However, when this bid price is low (as in the Nash equilibrium or rational expectation equilibrium of the type studied in Akerlof), increasing the price will seem attractive, precisely because the coarse categorization for ask prices above this bid price will incorrectly lead buyers to think they can gain a lot in quality by just increasing slightly the bid

[3] These take place at a variety of prices due to the heterogeneity of quality, and in all such matches, there is trade, since we assume, as does Akerlof, that there is always gain from trade.

price.[4] More specifically, buyers believe that quality within a category is constant, and equal to the empirical average quality within the category. Thus, to the buyers it will seem that raising the bid price slightly above p^* has a slight effect on the price to be paid, but increases quality more than slightly. Thus, it will seem as if the utility function of the buyers has a local optimum above p^*. Whether this local optimum also seems like a global optimum to the buyers depends on the level of the current price p^*. For p^* below or at the the Nash equilibrium price, it will in fact seem to the buyers that the best response is above p^*. For p^* sufficiently above the Nash equilibrium price it will instead seem to the buyers that the best response is at the Nash equilibrium price.

After developing our main insights in a simplified setting, we develop various robustness checks in order to strengthen our conclusion that categorical learning leads to cycles in decentralized trading mechanisms. We then explain why in the presence of centralized markets, no cycle should be expected, and we suggest that decentralized markets of the OTC type may be a way to boost volumes of trade when agents form their expectations based on categorical learning applied to data from past recent transactions.

Our paper is related to a number of approaches in which players may form erroneous expectations about some aspects of the interaction. In the context of adverse selection economies studied here, prominent such approaches include the cursed equilibrium (Eyster and Rabin 2005), in which buyers may fail to relate the strategy of the seller to the quality, and the behavioral equilibrium (Esponda 2008), in which buyers additionally only observe the quality when there is trade (similar to what we assume). These approaches, as well as the analogy-based expectation equilibrium, consider steady-state notions of equilibria, and as a result cannot capture the possibility of cycling. The main observation in Esponda is that, in contrast to what can arise in a cursed equilibrium, one will always obtain lowere volumes of trade in adverse selection economies when only the quality of the traded good is observed, as compared to the rational benchmark. What our analysis adds to this is that, if the coarseness of the understanding is adapted to the available data, then convergence to a steady-state is typically invalidated, and instead the economy will cycle, thereby generating more trade than in the rational benchmark.

Building on Esponda and Pouzo (2016) and Berk (1966) there is a growing literature on games with players who are Bayesian learners with misspecified models (see also Spiegler 2016). In this literature, beliefs take the form of a probability distribution over a set of parameters. The support of the belief is identified with a model, and if the support does not contain the true value of the parameter, the model is said to be misspecified.[5]

[4] On a technical note, the reason why, at the Nash equilibrium price, increasing the price is viewed as strictly beneficial is that under the coarse categorization, the effect of such a price increase would seem of second order while the coarse categorization leads to an upward boost of the estimate of an upward shift that is of first order.

[5] Perhaps the most straightforward interpretation of our theory (and ABEE in general) is that our agents are frequentists, but one can reinterpret them as being Bayesian with a suitable prior.

A few recent papers identify *cycles* of beliefs in the context of misspecified models. In Esponda, Pouzo and Yamamoto (2021) and Bohren and Hauser (2021) (see also Nyarko 1991), the evidence accumulated while taking a particular action may push beliefs in a direction that makes another action seem optimal, and once this new action is taken the data that are being generated induce a belief that makes the previous action seem optimal again. In Fudenberg Romanyuk, and Strack (2017) cycles may arise from the fact that the learner never ceases to perceive an information value of experimenting with another action. Note that in all of these papers, the learner sticks to a *fixed model* and in every period the learner selects a belief that *maximizes fit* with the data generated by current behavior (typically minimizing Kullback–Leibler divergence from the true distribution as the sample size goes to infinity). By contrast, in our case the learner's beliefs cycle because the learner switches *between models*. Moreover, the switch is not purely driven by likelihood considerations (choosing the best-fitting parameters), but also by *probability mass constraints* (i.e., the bias–variance trade-off).[6]

We hope that our finding may help usunderstand better some implications of OTC markets in financial contexts, in particular in relation to the boost in volumes of trade and cycling. Fluctuations in assets traded on OTC markets are documented by e.g. Bao et al. (2011) and Ivashchenko and Neklyudov (2018) among others. For a review of the large theoretical literature on OTC markets, that focuses on search-theoretic models, see Weill (2020). Maurin (2020) develops a liquidity-based theory of fluctuations in OTC markets with asymmetric information. Exogenous liquidity shocks, in the form of shocks to sellers' valuations, cause fluctuations in the price and average quality of traded assets. Price movements are driven by the fact that if future liquidity is high, forward-looking buyers expect to easily resell a low-quality asset and bid up its price.[7] A related model of cycles based on liquidity shocks is due to Asriyan et al. (2019). In our case seller valuations are entirely endogenously determined, and buyers are not forward-looking. Instead fluctuations in prices are due to boundedly rational expectation formation.

[6]A few papers study the choice between misspecified models, e.g., Cho and Kasa (2015) and Gagnon-Bartsch, Rabin, and Schwartzstein (2020). Fudenberg & Lanzani (2020) and He & Libgober (2020) perform evolutionary analyses of the stability of different misspecified models that the agents may consider. However, none of these papers obtain cycles.

[7]Asset owners face persistent idiosyncratic liquidity shocks in the form of shocks to their private value of the asset. The supply of high-quality assets will disproportionately be due to low-valuation sellers. Suppose that in the current period most owners of high-quality assets currently have a low valuation. Then the average quality of supplied assets is high and buyers are willing to pay more. In the next period, these buyers will own a large share of high-quality assets that they will not sell unless they face a liquidity shock. Hence the average quality of supplied assets is now lower, and buyers are willing to pay less. Thus the market becomes illiquid. The supply of high-quality assets eventually increases because of liquidity shocks.

2 Model

2.1 Market

We present our model in the context of an OTC market for loans.[8] Banks and other financial intermediaries act as sellers of loans (borrowers) to buyers of debt (lenders). The traded objects, the loans, are indivisible objects with random quality ω distributed on $\Omega = [0, 1]$ according to a continuous density function g, with cumulative G. One may think of the quality ω of a loan as being inversely related to the default probability, say one minus the default probability. The valuation of the seller/borrower coincides with the quality ω. The corresponding valuation of a buyer/lender is $v = \omega + b$, where $b \in (0, 1)$ represents gains from trade.[9]

We model OTC markets by considering one-to-one trading mechanisms between pairs consisting of one seller and one buyer drawn at random from their respective pools. In each pair, the seller and the buyer act simultaneously. The seller quotes an ask price $a(\omega)$ that depends on the quality ω that he privately observes. The buyer quotes a bid price p that depends on her information. The buyer is assumed to be uninformed of the quality ω with probability $1 - \gamma$ and informed of the quality with probability γ, where γ is assumed to be small. The market mechanism is such that if $p < a$ then there is no trade, and if $p \geq a$ then trade occurs at price p. Hence, if there is trade the buyer obtains utility $u(p) = v - p$, and the seller obtains utility p. If there is no trade, the seller gets ω and the buyer gets 0. This modeling of the trading mechanism allows us to simplify the analysis of the strategy of the seller, since setting the ask price equal to the quality $a(\omega) = \omega$ is a weakly dominant strategy for the seller, just as bidding one's own valuation is a weakly dominant strategy in the second-price auction. In the rest of the paper we assume that the seller employs his weakly dominant strategy.

We also assume that $b < (g(1))^{-1}$ and that G has the *monotone reversed hazard rate property* (or equivalently that G is strictly log-concave). That is, for all p,

$$\frac{\partial}{\partial p}\left(\frac{g(p)}{G(p)}\right) < 0.$$

Moreover, we assume the following smoothness condition. For all p,

$$|g'(p)| < g(p).$$

[8] The formalism can readily be reinterpreted in terms of any decentralized market for bilateral exchange where one side has private information.

[9] We may assume that gains from trade are a random variable B with mean $\mathbb{E}[B] = b$ and variance σ^2. Draws of B are independent and identically distributed across individuals and across values of ω. All results are unchanged. We also briefly discuss the case in which gains from trade take a multiplicative form, as in Akerlof's original modeling of the market for lemons.

While not essential for our main conclusion regarding the presence of price cycles, these extra assumptions will simplify the analysis by ensuring that there is a unique interior Nash equilibrium and will allow us to more completely analyze the cycling phenomenon.

2.1.1 Buyer/Lender and Seller/Borrower Behavior

As already mentioned, we assume that each seller/borrower follows the weakly dominant strategy that sets $a(\omega) = \omega$. We also restrict attention to pure strategies of the buyer/lender. When the buyer is informed of the quality ω (with probability γ), she bids[10] $p = \omega = a$. When the buyer is uninformed of the quality, she chooses a current bid p^* that will be determined by her current understanding of the strategy of the seller and how it relates to quality. The choice of p^*, which is an essential feature of our model, will be described below. The resulting density of transactions is[11]

$$h(s) = \begin{cases} g(s) & \text{if } s \leq p^* \\ \gamma g(s) & \text{if } s > p^* \end{cases}.$$

As already mentioned, we will think of γ as being small. The technical effect of the small fraction of perfectly informed buyers is to generate information about quality associated with prices that are different from what uninformed buyers currently perceive to be their optimal bid price. Our results are qualitatively the same if we instead assume that no buyer is perfectly informed but each buyer experiments with a small probability (see Section 4.1.6).

2.1.2 Dynamics

There is a population of buyers and a population of sellers. Each population has measure one. Time is discrete and in every period buyers and sellers are matched to interact with each other in the market. Data available at the end of period t are used by buyers in period $t+1$ to form beliefs about how the strategy of the seller relates to the quality of the loan. These beliefs are next used by uninformed buyers to adjust their choice of bid price in period $t+1$. Our preferred interpretation is that there is a new generation of buyers and sellers in every period rather than the same buyers and sellers interacting over many periods.

[10] We implicitly assume here that an informed buyer also knows the strategy of the seller. We can rationalize this by requiring not only that such an informed buyer know the quality ω, but also that the seller know it.

[11] Observe that the total mass of transactions $G(p^*) + (1-\gamma)\,G(p^*)$ is less than 1 since matches with uninformed buyers do not result in transactions whenever $\omega > p^*$. The presented density is not normalized by this mass of transactions.

2.1.3 Data

The buyers at time t have information about bid and ask prices from period $t-1$ as well as information about whether trade occurred. This allows them to form empirically based beliefs about the distribution of ask prices and the probability of trade as a function of bid price. Buyers also have information about quality in all the cases where trade took place in period $t-1$, but they do not have information about quality in the cases where trade did not occur. Based on these data, buyers form beliefs about the link between the quality ω and the ask prices by forming categories in the data set as we will explain shortly.

Up to now, our model has borne a strong resemblance to that of Esponda (2008). In a similar adverse selection environment, he also assumes that only the quality of traded goods is observed. Yet, as we will show below, our model differs from his in the modeling of the cognitive limitations of buyers, as well as in how categorical reasoning can give rise to cycling, which does not arise in his model.

Coming back to the OTC application, it seems highly plausible that there will be more information about those interactions in which transactions took place, than about those in which no transaction took place, as assumed in our model. By contrast, our assumption that ask and bid prices are accessible for all transactions may perhaps seem extreme in light of the relative opaqueness that is considered characteristic of OTC markets. In OTC markets, information about ask prices and quality in other matches may be limited. Our model emphasizes the impact of the opaqueness on the quality of the non-traded goods rather than on the transaction prices, but we believe our cycling phenomenon is likely to arise for more general specifications of what is observed about the transaction prices. From the perspective of our model, what matters is the relative abundance of information about transactions with an ask price below and above p^*, respectively.

2.2 Categorization and Beliefs

2.2.1 Categorization

Categorization is used by uninformed buyers to understand the statistical link between ask price and quality. Let p^* be the bid price chosen by uninformed buyers in the last period. Since the share γ of informed buyers is small, for ask prices less than or equal to p^* there are data on many transactions, while for ask prices above p^* there are data on only a few transactions (corresponding to informed buyers with $\omega > p^*$). Therefore, it is easy for buyers to predict quality as a function of ask prices less than or equal to p^*, but difficult to predict quality as a function of ask prices above p^*. In our main analysis we assume that buyers have a perfect understanding of how quality varies as a function of ask prices less than or equal to p^*. By contrast, in order to predict quality as a function of ask prices above p^* buyers bundle sets of ask prices into categories, thereby generating a coarse understanding of the mapping between ask prices and quality. Roughly, we require each category to contain the same share of the data set of transactions, and each category to consist of interval ask prices defined so that the mass of observed transactions within that range is above a target threshold, viewed as a primitive of the model. In each category, the observed average quality is computed from the data, and it is used to relate any ask price in this range to this average quality. In Section 4.1.7 we examine what happens when categories are used to predict quality both below and above p^*, and suggest that the same qualitative insights arise.[12]

A buyer's categorization C is a partition of $(p^*, 1]$ into $k \geq 1$ categories $C_1, C_2, ..., C_k$. It is defined based on a collection of cutoffs

$$p^* = c_0 < c_1 < c_2 < ... < c_{k-1} < c_k = 1.$$

We assume that each category includes its left boundary point:

$$\begin{aligned} C &= \{C_1, C_2, ..., C_k\} \\ &= \{(c_0, c_1], (c_1, c_2], (c_2, c_3], ..., (c_{k-2}, c_{k-1}], (c_{k-1}, c_k]\}. \end{aligned}$$

Let $i(p)$ denote the index of the category to which price p belongs. All buyers use the same categorization and make the same predictions.

2.2.2 Categorization Formation

Buyers form a categorization on the basis of the distribution of transactions. They use a heuristic that aims to split the mass of transactions into equally sized bins. The minimal probability mass in each category is $\bar{\kappa} \leq 1$; i.e., each category should satisfy $\bar{\kappa} \leq \int_{c_{i-1}}^{c_i} h(s)\, ds$, where $h(\cdot)$ denotes the density of transactions for the various qualities. Thus the number of categories is

$$k^+ = \max\left\{1, \left\lfloor \frac{1}{\bar{\kappa}} \int_{p^*}^{1} h(s)\, ds \right\rfloor \right\} = \max\left\{1, \left\lfloor \frac{\gamma(1 - G(p^*))}{\bar{\kappa}} \right\rfloor \right\},$$

each with a density of $\kappa^+ = \gamma(1 - G(p^*))/k^+$. Note that a categorization C is described completely by the parameters k^+, and p^*. Thus, taking the quality distribution G as

[12] The data are also used to estimate the probability of different ask prices. In principle, buyers may use a separate categorization for this purpose. However, since the total number of ask prices will be much larger than the number of ask prices associated with transactions, we simplify the analysis by assuming that buyers estimate the distribution of ask prices without the help of categorizations. Moreover, quality is probably best viewed as multidimensional and hence the modelling assumption that quality is unidimensional may obscure the fact that it is naturally more difficult to estimate the distribution of quality than the distribution of ask prices (even if the data sets are of the same size).

given, the categorization C is completely determined by[13] $\bar{\kappa}$, γ, and p^*. Appendix A.3 provides details on how the numbers of categories below and above p^* are determined as $\bar{\kappa} \to 0$ and $\gamma \to 0$.

2.2.3 Beliefs Based on Categories

Since the buyers and sellers constitute continuum populations with measure one, the empirical distribution of prices and qualities will (be assumed to) coincide with the true distributions.[14] For a bid price p less than or equal to p^* the buyers understand that expected quality in case of trade is $\mathbb{E}[\omega|\omega \leq p]$. Since they know that the probability of trade given a bid price p is $G(p)$ they correctly believe that the expected utility of bidding $p \leq p^*$ is $G(p)(\mathbb{E}[\omega|\omega \leq p] + b - p)$.

The buyers' expectation about quality in category C_i is derived from the actual quality of transactions associated with ask prices within category C_i. The expected value for a buyer from trades with an ask price in C_i is

$$v(C_i) = \mathbb{E}[\omega|c_{i-1} < \omega \leq c_i] + b,$$

and the expected quality of objects with an ask price in C_i is

$$\mathbb{E}[\omega|c_{i-1} < \omega \leq c_i] = \frac{\int_{c_{i-1}}^{c_i} \omega h(\omega)\, d\omega}{\int_{c_{i-1}}^{c_i} h(s)\, ds} = \frac{\int_{c_{i-1}}^{c_i} \omega g(\omega)\, d\omega}{\int_{c_{i-1}}^{c_i} g(s)\, ds}.$$

It follows that buyers form the following belief about the expected payoff from bidding p:

$$\pi^{CE}(p|p^*) = \begin{cases} G(p)(\mathbb{E}[\omega|\omega \leq p] + b - p) & \text{if } p \leq p^* \\ \\ G(p^*)(\mathbb{E}[\omega|\omega \leq p^*] + b - p) & \\ +\sum_{j=1}^{i(p)-1}(G(c_j) - G(c_{j-1}))(v(C_j) - p) & \text{if } p > p^*. \\ +\left(G(p) - G\left(c_{i(p)-1}\right)\right)\left(v\left(C_{i(p)}\right) - p\right) & \end{cases}$$

The reference to p^* reflects the fact that C is completely determined by p^* together with the parameters $\bar{\kappa}$ and γ.

As an illustration consider the case of a *uniform quality distribution* g, i.e., $G(p) = p$

[13] The minimum probability mass per category $\bar{\kappa}$ induces a minimum number of categories $\bar{k}$ as a function of $\bar{\kappa}$, p^*, and γ. One may consider a categorization heuristic that takes the total number of categories $\bar{k}$ as given and derives the minimum average probability mass per category $\bar{\kappa}$, as a function of $\bar{k}$, p^*, and γ. The results are qualitatively identical.

[14] On a literal reading, our continuum population assumption may suggest that buyers have access to data that allow them to use an arbitrarily fine-grained categorization also for ask prices above p^*, since even a fraction of a continuum population is also a continuum. However, the continuum assumption is merely a mathematical convenience that allows us to avoid the stochasticity that would arise in a model with finite populations. Working with large but finite populations would lead to complications but would not affect the main insights.

for all $p \in [0,1]$. Each category $C_i = [c_{i-1}, c_i]$ has the same width $c_i - c_{i-1} = \gamma(1-p^*)/k^+$ (where $k^+ = \max\{1, \lfloor \gamma(1-p^*)/\bar{\kappa} \rfloor\}$). The expected valued to a buyer from trades with an ask price in C_i is $v(C_i) = (c_i + c_{i-1})/2 + b$. Figure 1 illustrates this for parameter values $\gamma = 0.3$, $\bar{\kappa} = 0.1$, and $b = 0.3$ and four different values of p^*. The parameter assumptions imply that if $p^* \in \{0.2, 0.3\}$ there are two categories above p^*, whereas if $p^* \in \{0.4, 0.5\}$ there is only one category above p^*. Note that when $p^* = 0.2$ the buyer will perceive utility to be maximized by p somewhere between 0.3 and 0.4. However, when $p^* = 0.3$ it will seem to the buyer that utility is maximized at $p \approx 0.4$, when $p^* = 0.4$ it will seem that utility is maximized around $p \approx 0.5$, and when $p^* = 0.5$ it will seem that utility is maximized around $p \approx 0.3$.

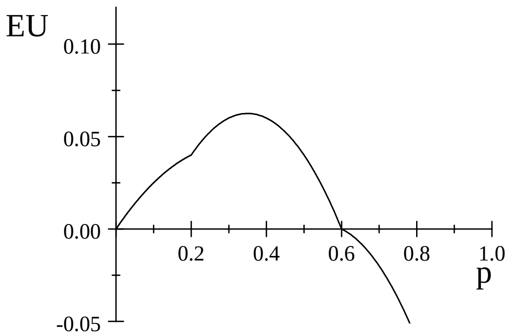

$p^* = 0.2$

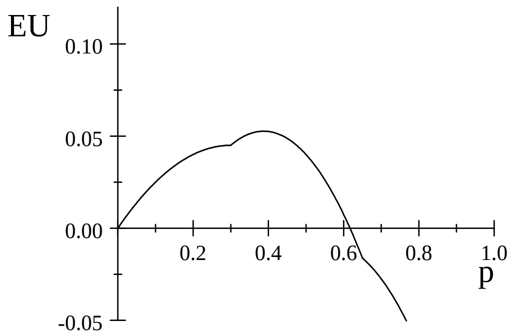

$p^* = 0.3$

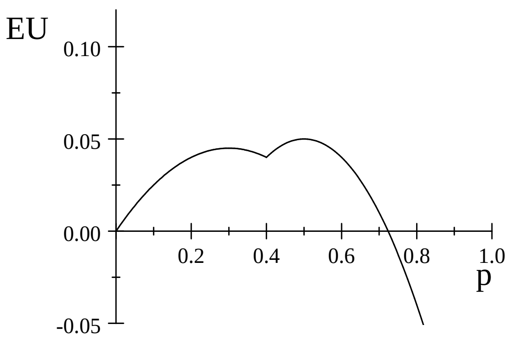

$p^* = 0.4$

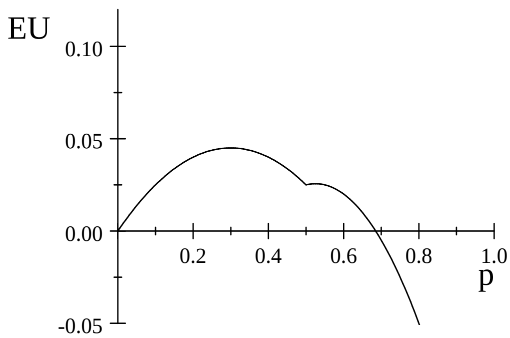

$p^* = 0.5$

Figure 1

2.2.4 Dynamics

Our dynamic system is completely characterized by the sequence of prices p_t^* quoted by uninformed buyers in period t. Given current price p_t^*, the next price p_{t+1}^* is chosen so that

$$p_{t+1}^* = \arg\max_p \pi^{CE}(p; p_t^*)$$

where $\pi^{CE}(\cdot;\cdot)$ is defined as in subsection 2.2.3.

We are interested in understanding the sequence p_t^*, and our main result will establish that this sequence must be cyclical no matter what the initial condition is.

3 Results

3.1 Nash Benchmark

In a Nash equilibrium buyers have correct expectations about the mapping between ask price and quality. They maximize

$$\pi^{NE}(p) = \int_{\omega=0}^{p} (\omega + b - p)\, g(\omega)\, d\omega = G(p)\left(\mathbb{E}[\omega|\omega \leq p] + b - p\right).$$

In the case of a uniform quality distribution g this becomes

$$\pi^{NE}(p) = p \cdot \left(\frac{p}{2} + b - p\right) = p\left(b - \frac{p}{2}\right),$$

and so the solution for buyers is $p^{NE} = b$. More generally, we have the following result:

Proposition 1 *There exists a unique Nash equilibrium in which the bid price p^{NE} of uninformed buyers is uniquely defined by*

$$\frac{g\left(p^{NE}\right)}{G\left(p^{NE}\right)} = \frac{1}{b}.$$

Proof. Note that

$$\begin{aligned}
\frac{\partial}{\partial p}\left(\mathbb{E}[\omega|\omega \leq p]\right) &= \frac{1}{G(p)} p g(p) - \left(\int_{\omega=0}^{p} \omega g(\omega)\, d\omega\right) \frac{g(p)}{G(p)^2} \\
&= \frac{g(p)}{G(p)}\left(p - \int_{\omega=0}^{p} \omega \frac{g(\omega)}{G(p)} d\omega\right) \\
&= \frac{g(p)}{G(p)}\left(p - \mathbb{E}[\omega|\omega \leq p]\right).
\end{aligned}$$

Thus

$$\begin{aligned}
\frac{\partial}{\partial p}\pi^{NE}(p) &= g(p)\left(\mathbb{E}\left[\omega|\omega \le p\right] + b - p\right) + G(p)\left(\frac{\partial}{\partial p}\left(\mathbb{E}\left[\omega|\omega \le p\right]\right) - 1\right) \\
&= g(p)\left(\mathbb{E}\left[\omega|\omega \le p\right] + b - p\right) + G(p)\left(\frac{g(p)}{G(p)}\left(p - \mathbb{E}\left[\omega|\omega \le p\right]\right) - 1\right) \\
&= g(p)\left(\mathbb{E}\left[\omega|\omega \le p\right] + b - p\right) + g(p)\left(p - \mathbb{E}\left[\omega|\omega \le p\right]\right) - G(p) \\
&= g(p)\, b - G(p),
\end{aligned}$$

and so the first-order condition of $\max_p \pi^{NE}(p)$ is

$$\frac{g(p)}{G(p)} = \frac{1}{b},$$

and the second-order condition is satisfied in virtue of the assumption that $|g'(p)| < g(p)$. Notice that $\lim_{p\to 0} \frac{g(p)}{G(p)} = \infty$ and $\frac{g(1)}{G(1)} = g(1)$. Hence, by the assumption that $g(1) < 1/b$ and $\frac{\partial}{\partial p}\left(\frac{g(p)}{G(p)}\right) < 0$, the first-order condition has a unique solution that is interior. ■

3.2 Categorical Learning

Our main result is that the sequence of p_t^* in the categorical learning model has no rest point and must cycle over finitely many values $p^{(k)}$, one of them being p^{NE} as previously characterized, and the others being above p^{NE}. In order to establish this, we first derive three properties related to how p_{t+1}^* varies with p_t^* depending on whether p_t^* is below, above, or equal to p^{NE}. These properties are referred to as lemmata and are proven in the Appendix.

The first property demonstrates that if at time t it is the case that $p_t^* = p^{NE}$, then at time $t+1$ uninformed buyers bid $p_{t+1}^* > p^{NE}$.

Lemma 1 *If $p_t^* = p^{NE}$ then*

$$p_{t+1}^* = \arg\max_{p\in[0,1]} \pi^{CE}\left(p|p^{NE}\right) > p^{NE}.$$

The next property demonstrates that if at time t it is the case that $p_t^* > p^{NE}$, then at time $t+1$ buyers either bid $p_{t+1}^* = p^{NE}$ or $p_{t+1}^* > p_t^*$.

Lemma 2 *If $p_t^* > p^{NE}$, then either*

$$p_{t+1}^* = \arg\max_{p\in[0,1]} \pi^{CE}\left(p|p_t^*\right) = p^{NE}$$

or

$$p_{t+1}^* = \arg\max_{p\in[0,1]} \pi^{CE}\left(p|p_t^*\right) > p_t^*.$$

The third property demonstrates that if at time t it is the case that $p_t^* < p^{NE}$, then at time $t+1$ buyers bid $p_{t+1}^* > p_t^*$.

Lemma 3 *If $p_t^* < p^{NE}$, then*

$$p_{t+1}^* = \arg\max_{p \in [0,1]} \pi^{CE}(p|p_t^*) > p_t^*.$$

Roughly, these three properties can be understood as follows. As already mentioned, categorical reasoning induces uninformed buyers to correctly infer that the quality corresponding to an ask price a below p^* is a. On the other hand, the coarse bundling for ask prices above p^* leads uninformed buyers to incorrectly infer that ask prices slightly above p^* are associated with an average quality that lies strictly above p^*. Thus, a buyer would choose a bid price strictly above p^* whenever $p^* \leq p^{NE}$ as she would incorrectly perceive a jump in quality when increasing slightly the bid price above p^* (and any bid price below p^* would rightly be perceived to be suboptimal). This is in essence the content of lemmata 3 and 1. By contrast, when $p^* > p^{NE}$, the best bid price below p^* is rightly perceived to be p^{NE} and the same logic leads the uninformed buyer to either choose p^{NE} or a bid price strictly above p^* with the aim of taking advantage of the jump in the perceived quality when the ask price lies above p^*.

The above properties immediately imply that the price dynamic has no rest point, i.e., there is no p_t^* such that

$$p_{t+1}^* = \arg\max_{p \in [0,1]} \pi^{CE}(p|p_t^*) = p_t^*.$$

To see this, assume by contradiction that p^* is a rest point. By Lemma 3, it cannot be that $p^* < p^{NE}$ since $p_t^* = p^* < p^{NE}$ would imply that $p_{t+1}^* > p_t^* = p^*$. By Lemma 1, it cannot be that $p^* = p^{NE}$ since $p_t^* = p^*$ would imply that $p_{t+1}^* > p^{NE}$. Finally, by Lemma 2, it cannot be that $p^* > p^{NE}$ since $p_t^* = p^*$ would imply either that $p_{t+1}^* > p_t^*$ or that $p_{t+1}^* = p^{NE}$ and thus $p_{t+1}^* \neq p_t^*$ (given that $p_t^* = p^* \neq p^{NE}$).

Even though there is no rest point, we will establish that there is a price cycle. In order to establish convergence to the price cycle we need one more lemma, which is again proven in the Appendix.

Lemma 4 *There is some $\delta > 0$ such that if $p^* \leq p^{NE}$ then $\mathbb{E}[\omega|\omega \in C_1] > p^* + \delta$.*

We can now state our main result to the effect that there is global convergence to a price cycle that consists of the Nash price and one or more prices above the Nash price.

Proposition 2 *There exists an increasing sequence $(p^{(1)}, ..., p^{(m)})$ with $m \geq 2$ and $p^{(1)} = p^{NE}$ such that if $p_t^* = p^{(i)}$ for $i \in \{1, ..., m-1\}$ then $p_{t+1}^* = p^{(i+1)}$, and if $p_t^* = p^{(m)}$ then $p_{t+1}^* = p^{(1)}$. Moreover, the dynamic converges to the set $\{(p^{(1)}, ..., p^{(m)})\}$ from any initial price $p_0 \in [0,1]$.*

Proof. Assume, to derive a contradiction, that the sequence p_t^* is monotonic. Lemmata 1–3 imply that

$$p_{t+1}^* = \arg\max_{p\in[0,1]} \pi^{CE}(p|p_t^*) > p_t^*$$

for all t. Since $p_t^* \leq 1$ for all t, it follows that $p_t^* \to \bar{p}$ for some $\bar{p} > p^{NE}$ as $t \to \infty$. (To see that there is a $\bar{p} > p^{NE}$ note that if $p_1^* \geq p^{NE}$ then $p_t^* \geq p^{NE}$ for all t.) This implies $\left|p_{t+1}^* - p_t^*\right| \to 0$, which, by continuity of $\pi^{CE}(p|p_t^*)$, implies $\left|\pi^{CE}\left(p_{t+1}^*|p_t^*\right) - \pi^{CE}(p_t^*|p_t^*)\right| \to 0$. Since $\pi^{CE}(p|p_t^*) = \pi^{NE}(p)$ for $p \in [0, p_t^*]$, we have $\left|\pi^{CE}\left(p_{t+1}^*|p_t^*\right) - \pi^{NE}(p_t^*)\right| \to 0$, and consequently $\pi^{CE}\left(p_{t+1}^*|p_t^*\right) \to \pi^{NE}(\bar{p})$. Since the Nash equilibrium p^{NE} is unique it holds that $\pi^{NE}\left(p^{NE}\right) > \pi^{NE}(\bar{p})$, and since $\pi^{CE}(p|p_t^*) = \pi^{NE}(p)$ for $p \in [0, p_t^*]$ we get

$$\pi^{CE}\left(p_{t+1}^*|p_t^*\right) \to \pi^{NE}(\bar{p}) < \pi^{NE}\left(p^{NE}\right) = \pi^{CE}\left(p^{NE}|p_t^*\right).$$

This is in contradiction to $p_{t+1}^* = \arg\max_{p\in[0,1]} \pi^{CE}(p|p_t^*)$. We conclude that the sequence p_t^* is not monotonic. Lemmata 1–3 imply that it must be cyclical, consisting of cycles with p^{NE} and one or more price above p^{NE}.

Note that the preceding argument can be used to show, that starting at $p_1^* \geq p^{NE}$ there is convergence to the cycle, from which there is no escape. To see this, suppose (to obtain a contradiction) that there is some $p_1^* > p^{NE}$ that does not belong to the cycle (i.e., $p_1^* \neq p^{(1)}$ for all $i \in \{1, ..., m\}$), from which there is no convergence to the cycle. This means that $p_{t+1}^* > p_t^*$ for all t and $p_t^* \to \bar{p}$ for some $\bar{p} \in \left[p^{NE}, p^{(m)}\right]$ as $t \to \infty$.

It remains to show that starting at $p_1^* < p^{NE}$ there is convergence to the set $\left[p^{NE}, 1\right]$. Consider $p_t^* < p^{NE}$. By Lemma 3 we know that $p_{t+1}^* = \max_p \pi^{CE}(p|p_t^*) > p_t^*$. For $p > p_t^*$ the FOC for $\max_{p\in C_{i(p)}} \pi^{CE}(p|p_t^*)$ is

$$g(p)\left(v\left(C_{i(p)}\right) - p\right) - G(p) = 0 \Longleftrightarrow p = v\left(C_{i(p)}\right) - \frac{G(p)}{g(p)}.$$

Suppose that

$$p_{t+1}^* = \max_p \pi^{CE}(p|p_t^*) \in C_1 = (c_0, c_1] = (p_t^*, c_1].$$

Then p_{t+1}^* satisfies

$$p_{t+1}^* \geq v(C_1) - \frac{G\left(p_{t+1}^*\right)}{g\left(p_{t+1}^*\right)},$$

holding with equality in the case $p_{t+1}^* < c_1$. Note that $\frac{G(p)}{g(p)}$ is increasing in p. Thus $p_{t+1}^* \in C_1$ and $p_{t+1}^* < p^{NE}$ imply that

$$p_{t+1}^* > v(C_1) - \frac{G\left(p^{NE}\right)}{g\left(p^{NE}\right)} = v(C_1) - b = \mathbb{E}[\omega|\omega \in C_1].$$

Lemma 4 further implies that $p_{t+1}^* - p_t^* > \delta > 0$. Above we assumed that $p_{t+1}^* = \max_p \pi^{CE}(p|p_t^*) \in C_1$. Clearly, if $p_{t+1}^* \notin C_1$ then $p_{t+1}^* > c_1 > p_t^* + \delta$. ■

The cycling result has interesting implications in light of the OTC application. It suggests that in the cycle there will be periods with higher volumes of trade (as arising when the bid price lies strictly above p^{NE}) followed by a period with a lower volume of trade that corresponds to the standard functioning of the adverse selection market (and that follows a logic similar to that developed in Akerlof's market for lemons).

4 Discussion

The purpose of this section is threefold. First, we consider several variants of the OTC model in which we suggest that the cycling phenomenon we identified in the main model is robust to a number of variations. Second, we consider a centralized trading mechanism and suggest there that no cycling should be expected. Finally, we position our study in relation to other approaches to bounded rationality in adverse selection markets, and highlight the novelty of the cycling phenomenon.

4.1 Variants

4.1.1 Multiplicative Gains from Trade

In the main model, gains from trade were assumed to be additive. Now we consider the possibility that gains from trade are instead multiplicative (as in Akerlof 1970). Specifically, suppose that the buyer's valuation is $v = \beta\omega$, where $\beta \in (1, 2)$ parameterizes the gains from trade. For simplicity, assume quality is uniformly distributed on $(0, 1)$. With multiplicative gains of trade the expected valued for a buyer from trades with an ask price in C_i is $v(C_i) = \mathbb{E}[\omega | c_{i-1} < \omega \leq c_i] \cdot \beta$. Let all other aspects of the model be unchanged. It turns out that the results are qualitatively similar to those obtained for the additive specification. (See Supplement S.2 for formal statements and proofs.) However, with a multiplicative specification adverse selection completely eliminates trade under the standard Bayes–Nash analysis: the unique Nash equilibrium is $p^{NE} = 0$. This observation is, of course, similar to that made in Akerlof. As in the additive specification, the Nash equilibrium is not a rest point under the dynamic induced by categorization-based learning: if in period t we are at the Nash price $p_t^* = p^{NE} = 0$, then in the next period buyers will perceive it to be utility-maximizing to state a bid price above the Nash price $p_{t+1}^* > 0$. For a learning model similar to the one studied with additive gains from trade, we can establish that categorization-based learning leads to cycle of the same kind as in the additive case. The cycle consists of the Nash price and at least one higher price. In the case where the cycle consists of more than two prices, the dynamic moves from one price to the next higher price until the highest price of the cycle is reached, from which there is a jump back to the lowest price, the Nash price.

4.1.2 When Some Buyers are Rational

In the main model, we have assumed that all buyers follow the categorical heuristic. Now suppose that the buyer population consists of two subpopulations: a fraction of individuals referred to as rational with a correct understanding of the mapping between ask prices and quality, and a complementary fraction of individuals using the categorical heuristic. Rational buyers will maximize π^{NE} and hence play p^{NE} regardless of what other players do. The question is whether they will affect the behavior of the non-rational players. When the non-rational buyers bid $p^* = p^{NE}$ the types are indistinguishable. Suppose the non-rational buyers bid $p^* > p^{NE}$. In this case, the data that are generated by the population as a whole will be different from the case without rational buyers. Let the fraction of rational buyers be μ. The total mass of transactions is then[15]

$$h(s) = \begin{cases} g(s) & \text{if } s \leq p^{NE} \\ (1-\mu)\, g(s) + \mu\gamma g(s) & \text{if } p^{NE} < s \leq p^* \\ \gamma g(s) & \text{if } p^* < s \end{cases}.$$

Regardless of the size of the fraction of rational buyers μ, there will be few transactions in the interval $(p^*, 1]$ and hence a coarse understanding of the relationship between quality and ask prices above p^*. If μ is small enough there will be sufficiently many transactions in the interval (p^{NE}, p^*) to allow for a perfect or near-perfect understanding of the relationship between quality and ask prices less than or equal to p^*. In this case, the above result will be robust to the introduction of rational buyers. By contrast, if μ is sufficiently large there will be only a few transactions in the interval (p^{NE}, p^*). Thus, it may be natural to assume a coarse understanding on the whole of $(p^{NE}, 1]$. In that case, the cycle will vanish. Instead, the rational players play p^{NE} and the non-rational players play p^* in every period. Thus, the cycling phenomenon requires that the share of non-rational buyers be not too small.

4.1.3 Heterogeneous Categorizations

Another relevant variation of the basic model is to allow different buyers to use different categorizations, such as arise from the use of different values of $\bar{\kappa}$. Consider the case of

[15] More explicitly, the distribution of transactions of Nash players is

$$h^{NE}(s) = \begin{cases} g(s) & \text{if } s \leq p^{NE} \\ \gamma g(s) & \text{if } s > p^{NE} \end{cases},$$

and the distribution of transactions of non-Nash players is, as before,

$$h^{NN}(s) = \begin{cases} g(s) & \text{if } s \leq p^* \\ \gamma g(s) & \text{if } s > p^* \end{cases}.$$

The total mass of transactions is then $h(s) = (1-\mu)\, h^{NN}(s) + \mu h^{NE}(s)$.

two kinds of buyers, denoted by A and B, who use different values of $\bar{\kappa}$. Suppose that all buyers initially bid the same price $p_t^* = p^{NE}$ and have a perfect understanding of quality for ask prices less than or equal to $p_t^* = p^{NE}$. The two different kinds of buyers may now have different beliefs about their best responses but they will both perceive their best responses to be above p_t^*. Without loss of generality we may assume that $p_t^* < p_{t+1}^{*A} < p_{t+1}^{*B}$. The behavior in period $t+1$ will generate many transactions in the interval $\left[0, p_{t+1}^{*A}\right]$, and few transactions in the interval $\left(p_{t+1}^{*B}, 1\right]$, and so in period $t+2$ we can assume a perfect understanding of quality for ask prices in $\left[0, p_{t+1}^{*A}\right]$ and a coarse understanding of quality for ask prices in $\left(p_{t+1}^{*B}, 1\right]$. Whether the understanding of quality for ask prices in $\left(p_{t+1}^{*A}, p_{t+1}^{*B}\right]$ is fine or coarse depends on whether the fraction of type-B buyers is large or small. In either case neither type will perceive it in their best interest to play the same bid price in period $t+2$ as in period $t+1$. As in the case of homogeneous categorizations it cannot be the case that $p_t^{*A} < p_{t+1}^{*A}$ and $p_t^{*B} < p_{t+1}^{*B}$ forever. Eventually either buyer type will switch back to p^{NE}. Thus, we will have cycling also in the case of heterogeneous categorizations, albeit of a more complicated kind than in the homogeneous case.[16]

4.1.4 Longer Memory

We have assumed that buyers base their categorizations and beliefs on data from the previous period only. What happens if buyers use data from the last $m \geq 2$ periods instead? In period t the highest period in the last m periods is $p_t^{*\max} = \max_{\tau \in \{t-m,\ldots,t-1\}} p_\tau^*$. The buyers in period t have a perfect understanding of how quality depends on ask prices less than or equal to $p_t^{*\max}$ (assuming that the amount of data needed for a perfect understanding is not larger in the case of $m \geq 2$ than in the case of $m = 1$). The coarseness of the categorization above $p_t^{*\max}$ depends on how much data have been generated above $p_t^{*\max}$ over the past m periods. If m is large enough (or rather if $m\gamma$ is large enough) then there might be enough data to afford a perfect or near-perfect understanding of quality above $p_t^{*\max}$. In that case, the best response is the Nash price for any t, and there is no cycle. If instead m is small enough (or rather if $m\gamma$ is small enough) then buyers will have a coarse understanding of quality above $p_t^{*\max}$. In this case there will be a cycle where p^{NE} is played for m periods before there are a number of periods with higher prices. To see this, note that if $p_t^{*\max} = p^{NE}$ then $p_t^* > p^{NE}$ and so in the next period $p_{t+1}^{*\max} = p_t^*$ and then $p_{t+1}^* > p_t^*$ or $p_{t+1}^* = p^{NE}$. To illustrate the new form of cycling, assume that $p_{t+1}^* = p^{NE}$; then in the next $m-1$ periods $p^{*\max} = p_t^*$ and $p^* = p^{NE}$, i.e., $p_j^{*\max} = p_t^*$ and $p_j^* = p^{NE}$ for $j \in \{t+1, t+2, ..., t+m\}$.

[16] One can view the situation studied in Section 4.1.3 as a limiting case of the one studied here in which the rational type would use $\bar{\kappa} = 0$.

4.1.5 Buyer Competition

Up to now, we have looked at one buyer–one seller interactions. We may alternatively consider matches consisting of more than one agent either on the buyer or on the seller side, in order to reflect that there is more competitive pressure on one side of the market. When there is more than one seller and just one buyer, everything boils down to considering the seller with the highest quality and so no qualitative changes should be expected in this case. When there is more than one buyer and just one seller, buyers engage in a kind of Bertrand competition. To simplify the analysis, we will assume that when several buyers are matched to a seller, they are either uninformed of the quality ω and make their choices of bid price based on categorical reasoning as described above, or they are informed of ω in which case they will now bid $\omega + b$ (instead of ω in the one seller–one buyer case) following the logic of Bertrand competition.

The more interesting aspect concerns the bidding of uninformed buyers. Each buyer is willing to outbid the opponent and increase her bid up to the point where she perceives that she earns zero expected profit. The perceived profit for buyer i bidding p_i when the other buyer bids p_{-i}, with categories and beliefs formed on the basis of price p_t^* as chosen by uninformed buyers in period t, is

$$\pi_i^B\left(p_i, p_{-i}|p_t^*\right) = \begin{cases} \pi_i^{CE}\left(p_i|p_t^*\right) & p_i > p_{-i} \\ \frac{1}{2}\pi_i^{CE}\left(p_i|p_t^*\right) & p_i = p_{-i} \\ 0 & p_i < p_{-i}. \end{cases}$$

In period $t+1$, uninformed buyers will choose p_{t+1}^* such that

$$\pi_i^{CE}\left(p_{t+1}^*, p_{t+1}^*|p_t^*\right) = 0.$$

A rest point of the dynamic, denoted by p^B, would have to satisfy $\pi_i^B\left(p^B, p^B|p^B\right) = 0$. In order for $p_t^* = p^B$ to be a rest point it must be the case that $\pi_i^B\left(p_i, p^B|p^B\right) \leq 0$ for all $p_i > p^B$. If this is not the case then the buyer competition will lead to a higher price $p_{t+1}^* > p_t^*$ in the next period. Later periods may lead to even higher prices but eventually, when in some period $t+\tau$ it holds that $\pi_i^B\left(p_i, p_{t+\tau}^*|p_{t+\tau}^*\right) \leq 0$ for all $p_i > p_{t+\tau}^*$ (and $\pi_i^B\left(p_{t+\tau}^*, p_{t+\tau}^*|p_{t+\tau}^*\right) < 0$), the dynamic will return to $p_t^* = p^B$.

For the uniform case, and assuming that there is always a single category above p^* (i.e., $\bar{\kappa}$ sufficiently large and γ sufficiently small), one can show that if $b < 1/8$ then there is a cycle, such that if $p_t^* = p^B = 2b$ then $p_{t+1}^* > \frac{1}{2} - 2b$. There may be more than two prices in the cycle (if $b > \frac{1}{16}$) but eventually the dynamic returns to p^B. (See Supplement S.3 for details.)

This analysis suggests that cycling may still emerge in the presence of richer specifications of the competitive pressures.

4.1.6 Experimentation or Noise

In the model presented above, observations of quality in transactions with quality above the equilibrium price p^* are generated by the fraction γ of matches in which the buyer is perfectly informed about quality. Without these transactions there would be no data on which the uninformed buyers could base their beliefs about the relationship between quality and ask prices above p^*. However, transactions with quality above the equilibrium price p^* could also be generated due to experimentation, or simple noise in the uninformed buyers' choice of bid price.

Suppose that with probability $1-\xi$ a buyer chooses a best response given her current beliefs. With probability ξ a buyer experiments and randomly chooses an action for the purpose of exploration according to a continuous density d with full support on $[0,1]$. Suppose that buyers currently find p^* optimal. Noting that conditional on experimenting, a good of quality s would be traded with probability $1-D(s)$, we obtain that the distribution of transactions is

$$h(s) = \begin{cases} (1-\xi)\, g(s) + \xi\,(1-D(s))\, g(s) & \text{if } s \leq p^* \\ \xi\,(1-D(s))\, g(s) & \text{if } s > p^*. \end{cases}$$

With this distribution of transactions, one can (qualitatively) recover the same results as in the main model.

4.1.7 Coarse Categorization below p^*

So far we have assumed that the buyers perfectly understand how quality is related to ask prices below p^*, but has a coarse understanding of this relationship for ask prices above p^*. We now analyze the case where a categorization is used also to predict quality for ask prices below p^*. The previous assumption that buyers perfectly understand quality below p^* may then be viewed as a limiting case of an infinitely fine-grained categorization below p^*. (See the Appendix for details.)

A buyer's categorization C is now a partition of $[0,1]$ into $k = k^- + k^+$ categories $C_1, C_2, ..., C_k$.. It is defined based on a collection of cutoffs

$$0 = c_{-k^-} < c_{-k^-+1} < ... < c_{-2} < c_{-1} < c_0 = p^* < c_1 < c_2 < ... < c_{k^+-1} < c_{k^+} = 1.$$

We assume that each category includes its left boundary point. The left-most category also includes its right boundary point:

$$\begin{aligned} C &= \{C_{-k^-}, , ...C_{-2}, C_{-1}, C_1, C_2, ..., C_{k^+}\} \\ &= \{(c_{-k^-}, c_{-k^-+1}], ..., (c_{-1}, c_0]\, (c_0, c_1], (c_1, c_2], ..., (c_{k-1}, c_k]\}. \end{aligned}$$

The numbers k^- and k^+ are set as the minimum ones so that each category has mass no less than $\bar{\kappa}$. Accordingly, we expect $k^-$to be large relative to k^+. More precisely, buyers form a categorization using a heuristic that aims to split the mass of transactions into equally sized bins where the minimal probability mass in each category is $\bar{\kappa} \leq 1$.

The buyers' expectation about quality in category C_i is derived from the actual quality of transactions associated with ask prices within category C_i as before. The buyers' belief about the expected payoff from bidding p is now

$$\pi^{CE}(p|p^*) = \sum_{j=1}^{i(p)-1} \left(G(c_j) - G(c_{j-1})\right)\left(v(C_j) - p\right) + \left(G(p) - G\left(c_{i(p)-1}\right)\right)\left(v\left(C_{i(p)}\right) - p\right).$$

In Figure 2 we illustrate this for parameter values $\gamma = 0.3$, $\bar{\kappa} = 0.1$, and $b = 0.3$ and four different values of p^*. These are the same values as in case of an arbitrarily fine-grained understanding of quality below p^* (i.e., the case of $k^- \to \infty$). Again, the parameter assumption implies that if $p^* \in \{0.2, 0.3\}$ there are two categories above p^*, whereas if $p^* \in \{0.4, 0.5\}$ there is only one category above p^*.

The picture is very similar to the case of a perfect understanding of quality below p^*. When $p^* = 0.2$ then the buyer will perceive utility to be maximized by p somewhere between 0.3 and 0.4. However, when $p^* = 0.3$ it will seem to the buyer that utility is maximized at $p \approx 0.4$, when $p^* = 0.4$ it will seem that utility is maximized around $p \approx 0.5$, and when $p^* = 0.5$ then it will semme that utility is maximized around $p \approx 0.3$.

Assuming a uniform quality distribution g we are able to obtain a number of analytical results. There is convergence to an interval containing the Nash price b. Starting from any initial p_0^*, as $t \to \infty$ the dynamic will converge to, and never escape, the set

$$P^*(\bar{\kappa}, \gamma) = \left(b - \frac{1}{2\left(\frac{1}{\bar{\kappa}} - 1\right)}, b + \frac{1}{2\left(\frac{\gamma}{\bar{\kappa}} - 1\right)}\right).$$

As the minimal mass of categories $\bar{\kappa}$ decreases and approaches 0 the set $P^*(\bar{\kappa}, \gamma)$ converges to (the singleton set consisting of) b. One might conjecture that $p^* = b$ is a steady-state for a small enough value of $\bar{\kappa}$. However, this is not the case. If $p^* = b$ then the buyer's perceived best response is bounded away from b. In the special cases of $k^- = 1$ and $k^+ = 1$ we can identify a cycle with two price levels. (For formal statements and proofs see Supplement S.1.)

We expect our main result, regarding convergence to a price cycle, to apply in this more elaborate setting when there are sufficiently many categories below p^* and sufficiently few categories above p^*. We examine this numerically by examining how price movements vary when we vary the information parameter γ, and the categorization coarseness parameter $\bar{\kappa}$. (The results are presented in Supplement S.1.2.) In general, increasing the share of informed buyers γ, and decreasing the required mass per category

$\bar{\kappa}$, dampens fluctuations by increasing the number of categories above p^*.

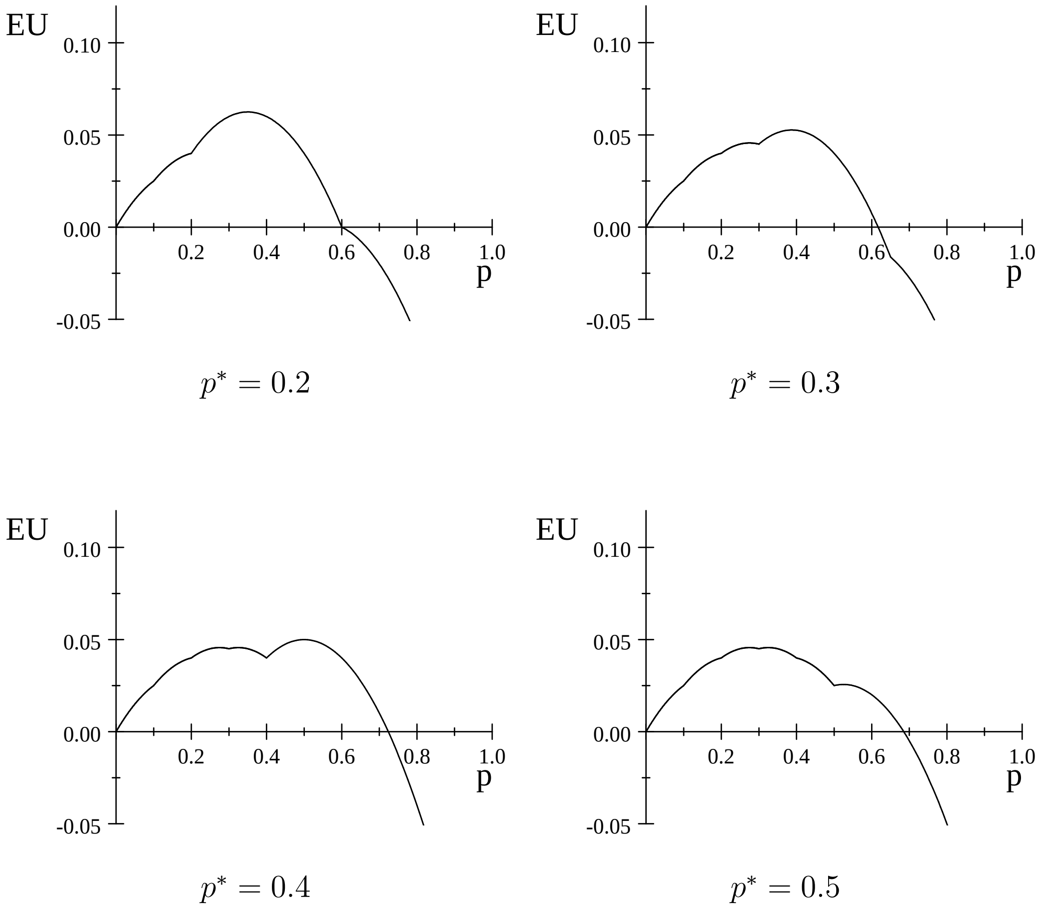

Figure 2

4.2 Centralized Trading Mechanism

The market studied so far was decentralized in the spirit of OTC markets with trades involving small sets of agents (one seller–one buyer in the basic formulation). In that context we found that when buyers rely on the categorical heuristic, prices in the corresponding dynamic model cycle and never fall short of the price that would arise if buyers had a perfect understanding of the relationship between ask prices and quality (as in the standard analysis).

It is of interest to analyze what would happen in similar observational contexts if instead of considering decentralized trading procedures, the trading mechanism was centralized with a single market governing all transactions. We will illustrate that in such a case, no cycling should be expected and (at least in some cases) the induced volume of trade may be smaller than that arising in decentralized markets.

To establish this, consider the following trading mechanism. In each period, all sellers and all buyers simultaneously submit their ask and bid prices to a central clearinghouse that chooses p^* to equate supply and demand. As implied by the centralized nature of the market, the traded goods are assigned randomly to those buyers who engage in trade. If we had the same number n of buyers and sellers and the clearing price was chosen to be the largest one that allows the clearinghouse to equate demand and supply, the sellers would have a weakly dominant strategy consisting of setting an ask price $a(\omega) = \omega$ when the quality is[17] ω. When considering the limit case with a continuum of buyers and sellers of equal mass, we will suppose that a seller sets $a(\omega) = \omega$ when selling a good of quality ω. We note that in the context of centralized trading mechanisms, it is irrelevant whether a buyer observes the quality of the sellers' goods to the extent that she has no control over which good she will receive (since the traded goods are assigned randomly). The average quality of goods traded at price p_t^* at time t is then $\mathbb{E}\left[\omega|\omega \leq p_t^*\right]$.

In line with our assumptions on decentralized markets, we assume that buyers observe the average quality of the goods sold as well as the distribution of ask prices in the previous period. Thus, at time $t+1$ a buyer's perceived expected value of an object is[18] $\mathbb{E}\left[\omega|\omega \leq p_t^*\right] + b$. This implies that at time $t+1$ a buyer will not post a price $p_{t+1} > \mathbb{E}\left[\omega|\omega \leq p_t^*\right]+b$, since this would be dominated by posting a price equal to $\mathbb{E}\left[\omega|\omega \leq p_t^*\right]+b$. Moreover, in the limit of an infinite buyer population, the probability of a single buyer affecting the equilibrium price is zero. Hence, posting a price $p_{t+1} < \mathbb{E}\left[\omega|\omega \leq p_t^*\right] + b$ is weakly dominated by setting[19] $p_{t+1} = \mathbb{E}\left[\omega|\omega \leq p_t^*\right] + b$. Consequently, the market clearinghouse (Walrasian auctioneer) would have to set a price $p_{t+1}^* = \mathbb{E}\left[\omega|\omega \leq p_t^*\right] + b$ to equate demand and supply, since at any price $p < \mathbb{E}\left[\omega|\omega \leq p_t^*\right] + b$ there is excess demand, and at any price $p > \mathbb{E}\left[\omega|\omega \leq p_t^*\right] + b$ there is excess supply.

A steady-state price p^C of the centralized market satisfies $\mathbb{E}\left[\omega|\omega \leq p^C\right] + b = p^C$. It is readily verified that if $b < 1 - \mathbb{E}\left[\omega\right]$ (or $b < 1/2$ in the uniform case) the steady-state price p^C would need to be interior.[20] We note that p^C coincides with p^B as described above in the decentralized case when there are at least two buyers matched with a seller and these are rational. (This should come as no surprise given that if the masses of buyers and sellers are the same, and if the price is interior, (the effective competitive pressure is on the buyers' side.) It is also the rational expectation equilibrium of the type studied in Akerlof.

[17] This follows the same logic as in the one-object second-price auction and can be viewed as an illustration of Vickrey's generalized version of it.

[18] In the case of $p_t^* = 0$ there is no trade and we assume that $\mathbb{E}\left[\omega|\omega \leq p_t^*\right] = \lim_{p_t^* \to 0} \mathbb{E}\left[\omega|\omega \leq p_t^*\right] = 0$.

[19] Alternatively, we may replace the simultaneous bid assumption with a gradual process in which the auctioneer starts with a low price and then gradually increases it. Buyers once they drop out cannot come back. Sellers once they join cannot drop out. The price increase stops when there is an equal number of buyers and sellers.

[20] When $b < 1 - \mathbb{E}\left[\omega\right]$ a steady price $p^C = 1$ can be ruled out as it means that all sellers would want to sell, but the buyer's perceived utility is $\mathbb{E}\left[\omega|\omega \leq p^C\right] + b - p^C = \mathbb{E}\left[\omega\right] + b - 1 < 0$, thereby leading to a contradiction. A steady-state price $p^C = 0$ can be ruled out since it leads to zero supply but positive demand due to $E\left[\omega|\omega \leq p^C\right] + b - p^C = b > 0$.

We are interested to know whether the dynamic $p^*_{t+1} = \mathbb{E}[\omega|\omega \leq p^*_t] + b$, which would arise from the learning dynamic in the centralized trading case, converges to the steady-state p^C. It turns out that under our maintained assumption that G is strictly log-concave p^C is locally attracting. That is, there is an open interval $(\underline{p}, \overline{p})$ around p^C, such that if $p^*_0 \in (\underline{p}, \overline{p})$ then $p^*_t \to p^C$. Furthermore, if we impose the additional condition that

$$\frac{g(p) \int_0^p G(\omega)\, d\omega}{G(p)^2} < 1, \tag{1}$$

then p^C is globally attracting. (See the Appendix for details.)

In order to compare the volume of trade in the decentralized and centralized cases it may be useful to assume that ω is uniformly distributed on $(0, 1)$. The condition for a steady-state in the centralized market case becomes $\frac{1}{2}p^C + b = p^C$, and so $p^* = 2b$. Since only a fraction $2b$ of sellers have a valuation less than $2b$, this value it is also the volume of trade. By comparison, the Nash equilibrium in the decentralized (one seller–one buyer) case is $p^{NE} = b$, and in our learning dynamic the prices take values above p^{NE}. When b is small enough, the expected price arising in the decentralized case is larger than p^*, showing that there is more trade under the decentralized trading mechanism.

To be more specific, in the case of a uniform distribution one can show (Supplement S.3, Proposition S6) that if there is always a single category above p^* then the cycle consists of exactly two prices, the Nash price $p^{NE} = b$ and a higher price equal to $(1 + 3b)/4$. The cycle alternates between these two prices and so the average price is $(1 + 7b)/8$. Since there is always trade in matches involving informed buyers, the resulting average trade volume is at least $(1 + 7b)/8$. Thus, average trade volume in the decentralized case exceeds that of the centralized case whenever $(1 + 7b)/8 > 2b$ or equivalently $b < 1/9$, i.e., gains from trade of less than 11%.

For larger values of b, there may be more expected trade in the centralized trading mechanism than in our main decentralized version. However, the conclusion that there is more trade in the decentralized than in the decentralized case would always hold if in the decentralized case we were to assume that there is more than one buyer in every match involving a seller (since then $p^B = p^*$ and we may have cycles involving prices above p^B in the decentralized case). The conclusion that there is more trade in the centralized than in the decentralized case would also hold when gains from trade take a multiplicative form, since in this case (as in Akerlof) there would be no trade at all in the centralized trading mechanism.

Altogether, these insights support the conclusion that decentralized trading mechanisms may lead to larger volumes of trade than centralized trading mechanisms, when buyers follow a categorical learning heuristic that is based on the available data from past recent trades.

4.3 Related Literature

The assumption that information about quality is only available for past instances of trade, together with the resulting lack of data above the current market price, is taken from Esponda (2008). He assumes that buyers are able to form an accurate belief about how quality depends on ask prices below the equilibrium price, but not necessarily on those above it. He considers a naive and a sophisticated alternative. Naive buyers believe that the average quality above the equilibrium price is the same as the average quality below the equilibrium price. Thus, naive buyers do not understand that increasing the bid price leads to transactions with higher quality. By contrast, sophisticated buyers do to some extent understand the positive relationship between ask prices and quality. If all buyers are naive the resulting behavioral equilibrium price is even lower than the Nash equilibrium price: adverse selection is aggravated by the naive buyers' pessimistic view of quality associated with prices above the equilibrium price. If instead buyers are sophisticated the resulting behavioral equilibrium price is bounded from below by the naive equilibrium price and bounded from above by the Nash equilibrium price (under the belief consistency conditions imposed by Esponda).

Eyster and Rabin (2005), see also Miettinen (2009), define a notion of cursedness, such that in our setting a fully cursed buyer fails entirely to take selection into account and believes that quality is independent of the ask price. Presumably such buyers entirely ignore the correlation between ask price and correlation that they can observe below the equilibrium price. In the case where all buyers are fully cursed the cursed equilibrium may be below or above the Nash equilibrium price depending on the size of gains from trade. In our setting the fully cursed price is below the Nash price if gains from trade are large ($b > 1/2$ in the uniform case) and above if gains from trade are small ($b < 1/2$ in the uniform case). A partially cursed buyer partially takes selection into account in the sense that her expected utility is a convex combination of what a cursed buyer and a Nash buyer would believe. Depending on the degree of cursedness the cursed equilibrium is somewhere between the Nash price and the price in an equilibrium with only fully cursed buyers.

Suppose that, as in Esponda's setup, buyers observe quality only in cases where trade took place, and suppose that, as in our model above, with some small probability a buyer is perfectly informed about quality. In the analogy-based expectations equilibrium (ABEE), due to Jehiel (2005) and Jehiel and Koessler (2008), buyers use categories to judge quality as a function of prices (as in our model) but categories are fixed. For simplicity, assume that all categories have the same width.[21] In such a setup one can

[21] Miettinen (2012) provides a refinement of ABEE, called a payoff-confirming analogy equilibrium (PAE), that requires that the payoffs that players obtain are consistent with their ABEE-induced beliefs about the distribution of play and types. In the lemons market that we consider, only the arbitrarily fine-grained categorization constitutes a PAE.

show that as the number of categories goes to infinity the ABEE price approaches the Nash equilibrium price, while if there is a single category then the ABEE price coincides with the fully cursed equilibrium price. A detailed comparison of behavioral equilibrium, cursed equilibrium, and ABEE is provided in Supplement S.5.

In summary, in the behavioral equilibrium prices are bounded from above by the Nash price, while in our case the price cycle is bounded from below by the Nash price. In the cursed equilibrium and the ABEE price may be above or below the Nash price but there is no cycling. The key facet of our model that generates cycles is the endogeneity of the categories: the empirical observations endogenously determine what model the agents' use to form their beliefs.

The kind of market that we have examined theoretically was implemented experimentally by Fudenberg and Peysakhovich (2016), for the case of a uniform quality distribution g. They compare treatments corresponding to $b = 0.3$ and $b = 0.6$ as well as different information conditions. They find prices above the Nash equilibrium, and they find that a lower value of b yields relatively more overshooting. This is in line with our theory. As mentioned above, we can show that if there is always a single category above p^* then the cycle consists of exactly two prices (Supplement S.3, Proposition S6). For $b = 0.3$ our then theory predicts a cycle between 0.3 and 0.475, where the latter price is 58% above the former price. For $b = 0.6$ our theory predicts a cycle between 0.3 and 0.7, where the latter price is 14% above the former price. Fudenberg and Peysakhovich find that buyers continue to bid above b even when provided with information about the distribution of quality and when they can observe quality associated with ask prices that did not result in trade. Providing information about average quality in the data set does bring bids down to an average of 4.12 (for $b = 0.3$).

5 Conclusion

Our analysis shows that decentralized trading mechanisms, in contrast to centralized ones, may generate cyclical patterns with volumes of trades never below that of the rational expectation benchmark. Specifically, this holds when uninformed traders rely on categorical learning to form expectations, and quality is observed only for traded assets. The identified mechanism provides a novel explanation of fluctuations in OTC markets. It also suggests a rationale for using OTC markets to generate higher volumes of trade as opposed to more traditional centralized markets, thereby explaining the widespread use of OTC markets despite improved information technology (Weill 2020; Riggs et al. 2020). Further empirical work is needed to assess the extent to which fluctuations in OTC markets are better explained by our mechanism or by more traditional ones in the presence of rational forward-looking agents. Further theoretical work may investigate the effect of market entry and exit during booms and busts.

6 References

Agarwal, S., Driscoll, J. C., Gabaix, X., and Laibson, D. (2008) Learning in the credit card market (National Bureau of Economic Research, Working Paper 13822.

Akerlof, G. A. (1970). The market for lemons: Quality and the market mechanism. *Quarterly. Journal Economics*, 84, 488–500.

Asriyan, V., Fuchs, W., and Green, B. (2019). Liquidity sentiments. *American Economic Review*, 109(11), 3813–48.

Bagnoli, M., and Bergstrom, T. (2005). Log-concave probability and its applications. *Economic Theory*, 26(2), 445–469.

Bao, J., Pan, J., and Wang, J. (2011). The illiquidity of corporate bonds. *Journal of Finance*, 66(3), 911–946.

Berk, Robert (1966). Limiting behavior of posterior distributions when the model is incorrect. *Annals of Mathematical Statistics*, 37, 51–58.

Bohren, J. A., and Hauser, D. N. (2021). Learning with heterogeneous misspecified models: characterization and robustness. Forthcoming in *Econometrica*.
Available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3783840

Erev, I., Haruvy, E. (2014). Learning and the economics of small decisions. In *The Handbook of Experimental Economics*, Vol 2, eds. Roth, A. E., Kagel, J. Princeton NJ: University Press.

Esponda, I. (2008). Behavioral equilibrium in economies with adverse selection" *American Economic Review* 98, 1269–91.

Esponda, I., and Pouzo, D. (2016). Berk–Nash equilibrium: A framework for modeling agents with misspecified models. *Econometrica*, 84(3), 1093–1130.

Esponda, I., Pouzo, D., and Yamamoto, Y. (2021). Asymptotic behavior of Bayesian learners with misspecified models. *Journal of Economic Theory*, 195, 1052–1060.

Eyster, E., and Rabin, M. (2005). Cursed Equilibrium. *Econometrica*, 73(5), 1623–72.

Fudenberg, D., and Peysakhovich, A. (2016). Recency, records, and recaps: Learning and nonequilibrium behavior in a simple decision problem. *ACM Transactions on Economics and Computation*, 4(4), 23.

Fudenberg, D., Romanyuk, G., and Strack, P. (2017). Active learning with a misspecified prior. *Theoretical Economics*, 12(3), 1155–1189.

Fudenberg, D., and Lanzani, G. (2020). Which misperceptions persist?. Working Paper.
Available at SSRN. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3709932

Gagnon-Bartsch, T., Rabin, M., and Schwartzstein, J. (2018). Channeled attention and stable errors. Working Paper.

He, K., and Libgober, J. (2020). Evolutionarily stable (mis)specifications: Theory and applications. arXiv preprint arXiv:2012.15007.

Ivashchenko, A., and Neklyudov, A. (2018). (In)frequently traded corporate bonds. Available at SSRN 3306124.

Jehiel, P. (2005). Analogy-based expectation equilibrium, *Journal of Economic Theory,* 123(2), 81–104.

Jehiel, P., and Koessler, F. (2008). Revisiting games of incomplete information with analogy-based expectations equilibrium, *Games and Economic Behavior*, 62(2), 533–57.

Maurin, V. (2020). Liquidity fluctuations in over-the-counter markets. Forthcoming in *Journal of Finance.*

Miettinen, T. (2009). The partially cursed and the analogy-based expectation equilibrium. *Economics Letters*, 105(2), 162–164.

Miettinen, T. (2012). Paying attention to payoffs in analogy-based learning. *Economic Theory*, 50(1), 193–222.

Mohlin, E. (2014). Optimal categorization. *Journal of Economic Theory,* 152, 356–381.

Nyarko, Y. (1991). Learning in mis-specified models and the possibility of cycles. *Journal of Economic Theory*, 55, 416–427.

Spiegler, R. (2016). Bayesian networks and boundedly rational expectations. *Quarterly Journal of Economics*, 131(3), 1243–1290.

Riggs, L., Onur, E., Reiffen, D., and Zhu, H. (2020). Swap trading after Dodd–Frank: Evidence from index CDS. *Journal of Financial Economics*, 137(3), 857–886.

Weill, P. O. (2020). The search theory of OTC markets. National Bureau of Economic Research Working paper 27354.

Appendix

A.1 Preliminaries

Note that $\lim_{p\uparrow c_i} \pi^{CE}(p|p^*) = \lim_{p\downarrow c_i} \pi^{CE}(p|p^*)$, for all $i \in \{1, ..., k-1\}$, implying that $\pi^{CE}(p|p^*)$ is continuous everywhere. Moreover, $\pi^{CE}(p|p^*)$ is piecewise differentiable with points of non-differentiability only at category boundaries. Consider the problem of maximizing $\pi^{CE}(p|p^*)$ under the restriction that the price $p > p^*$ belongs to category C_i,

$$\max_{p\in C_i} \pi^{CE}(p|p^*) = G(p^*)(\mathbb{E}[\omega|\omega \le p^*] + b) + \sum_{j=1}^{i(p)-1} (G(c_j) - G(c_{j-1}))(v(C_j)) \\ - G\left(c_{i(p)-1}\right)\left(v\left(C_{i(p)}\right)\right) + \max_{p\in C_i}\left(G(p)\left(v\left(C_{i(p)}\right) - p\right)\right).$$

The first derivative at $p \in \left(c_{i(p)-1}, c_{i(p)}\right)$ is

$$\frac{\partial \pi^{CE}(p|p^*)}{\partial p} = g(p)\left(v\left(C_{i(p)}\right) - p\right) - G(p).$$

The second derivative is

$$\begin{aligned} \frac{\partial^2 \pi^{CE}(p|p^*)}{\partial p^2} &= g'(p)\left(v\left(C_{i(p)}\right) - p\right) - 2g(p) \\ &= g(p)\left(\frac{g'(p)}{g(p)}\left(\mathbb{E}\left[\omega|\omega \in C_{i(p)}\right] + b - p\right) - 2\right) \\ &\le g(p)\left(\frac{|g'(p)|}{g(p)}(1+b) - 2\right) \\ &\le 2g(p)\left(\frac{|g'(p)|}{g(p)} - 1\right) < 0, \end{aligned}$$

where the inequality is verified using the assumption that $|g'(p)| < g(p)$. Thus $\pi^{CE}(p|p^*)$ is concave on each interval $\left(c_{i(p)-1}, c_{i(p)}\right)$.

A.2 Proof of Lemmata Needed for the Proof of Proposition 2

Proof of Lemma 1. Since $\pi^{CE}\left(p|p^{NE}\right)$ coincides with $\pi^{NE}(p)$ on $[0, p^*] = \left[0, p^{NE}\right]$, the constrained optimal $p \in [0, p^*]$ is at $p = p^* = p^{NE}$. Differentiating π^{CE} at $p \in C_1 =$

$\left(p^{NE}, c_1\right]$, and letting p go to p^{NE}, we obtain

$$
\begin{aligned}
\left.\frac{\partial \pi^{CE}\left(p|p^{NE}\right)}{\partial p}\right|_{p\downarrow p^{NE}} &= g\left(p^{NE}\right)\left(\mathbb{E}\left[\omega|\omega\in C_1\right]+b-p^{NE}\right)-G\left(p^{NE}\right) \\
&= G\left(p^{NE}\right)\left(\frac{g\left(p^{NE}\right)}{G\left(p^{NE}\right)}\left(\mathbb{E}\left[\omega|\omega\in C_1\right]+b-p^{NE}\right)-1\right) \\
&= G\left(p^{NE}\right)\left(\frac{1}{b}\left(\mathbb{E}\left[\omega|\omega\in C_1\right]+b-p^{NE}\right)-1\right) \\
&= \frac{G\left(p^{NE}\right)}{b}\left(\mathbb{E}\left[\omega|\omega\in C_1\right]-p^{NE}\right) \\
&= g\left(p^{NE}\right)\left(\mathbb{E}\left[\omega|\omega\in C_1\right]-p^{NE}\right)>0.
\end{aligned}
$$

Here, the third and fifth equalities use the fact that $g\left(p^{NE}\right)/G\left(p^{NE}\right)=1/b$. Since $\pi^{NE}(p)$ is continuous, the desired result is implied. ■

Proof of Lemma 2. Since $\pi^{CE}(p|p_t^*)$ coincides with $\pi^{NE}(p)$ on $[0,p_t^*]$, the constrained optimal $p\in[0,p_t^*]$ is at $p=p^{NE}<p_t^*$. Suppose that $\arg\max_{p\in[p_t^*,1]}\pi^{CE}(p|p_t^*)=p_t^*$ (requiring $\left.\frac{\partial\pi^{CE}(p|p_t^*)}{\partial p}\right|_{p\downarrow p_t^*}\leq 0$). Then by continuity of $\pi^{CE}(p|p_t^*)$, $\arg\max_{p\in[0,1]}\pi^{CE}(p|p_t^*)=p^{NE}<p_t^*$. ■

Proof of Lemma 3. Suppose, $p_t^*<p^{NE}$. Then the constrained optimal $p\in[0,p_t^*]$ is at p_t^* and

$$
\begin{aligned}
\left.\frac{\partial \pi^{CE}\left(p|p_t^*\right)}{\partial p}\right|_{p\downarrow p_t^*} &= \left. g(p)\left(v\left(C_{i(p)}\right)-p\right)-G(p)\right|_{p=p_t^*} \\
&= g(p_t^*)\left(\mathbb{E}\left[\omega|c_{i-1}<\omega\leq c_i\right]+b-p_t^*\right)-G(p_t^*) \\
&= G(p_t^*)\left(\frac{g(p_t^*)}{G(p_t^*)}\left(\mathbb{E}\left[\omega|c_{i-1}<\omega\leq c_i\right]+b-p_t^*\right)-1\right) \\
&\geq G(p_t^*)\left(\frac{g\left(p^{NE}\right)}{G\left(p^{NE}\right)}\left(\mathbb{E}\left[\omega|c_{i-1}<\omega\leq c_i\right]+b-p_t^*\right)-1\right) \\
&= G(p_t^*)\left(\frac{1}{b}\left(\mathbb{E}\left[\omega|c_{i-1}<\omega\leq c_i\right]+b-p_t^*\right)-1\right) \\
&= g(p_t^*)\left(\mathbb{E}\left[\omega|c_{i-1}<\omega\leq c_i\right]-p_t^*\right)>0.
\end{aligned}
$$

By continuity of π^{CE} we have $\arg\max_{p\in[0,1]}\pi^{CE}(p|p_t^*)>p_t^*$. ■

Proof of Lemma 4. Let $g^{\min}=\min_{\omega\in[0,1]}g(\omega)$ and $g^{\max}=\max_{\omega\in[0,1]}g(\omega)$. By the full-support assumption $g^{\min}>0$. Similarly, define $h^{\min}$ and $h^{\min}$, and note that

$h^{\min} = \gamma g^{\min}$ and $h^{\max} = \gamma g^{\max}$. Moreover, note that

$$\kappa^{+} \leq \int_{\omega \in C_1} h(\omega)\, d\omega = \gamma \int_{\omega \in C_1} g(\omega)\, d\omega \leq \gamma \int_{\omega \in C_1} g^{\max} d\omega = \gamma (c_1 - c_0)\, g^{\max}$$
$$\Rightarrow c_1 - c_0 \geq \frac{\kappa^{+}}{\gamma g^{\max}}.$$

We have

$$\int_{\omega \in C_1} h^{\min} d\omega = (c_1 - c_0)\, h^{\min} = (c_1 - c_0)\, \gamma g^{\min},$$

and

$$\int_{\omega \in C_1} \omega \frac{h^{\min}}{\int_{s \in C_1} h^{\min} ds} d\omega = \frac{c_1 + c_0}{2} = c_0 + \frac{c_1 - c_0}{2}.$$

We can use this to find a lower bound on $\mathbb{E}[\omega | \omega \in C_1]$:

$$\begin{aligned}
\mathbb{E}[\omega | \omega \in C_1] &\geq \frac{1}{\kappa^{+}} \left(\left(\kappa^{+} - \int_{\omega \in C_1} h^{\min} d\omega \right) c_0 + \left(\int_{\omega \in C_1} h^{\min} d\omega \right) \left(\frac{\int_{\omega \in C_1} \omega h^{\min} d\omega}{\int_{s \in C_1} h^{\min} ds} \right) \right) \\
&= \frac{1}{\kappa^{+}} \left(\left(\kappa^{+} - \int_{\omega \in C_1} h^{\min} d\omega \right) c_0 + \left(\int_{\omega \in C_1} h^{\min} d\omega \right) \left(c_0 + \frac{c_1 - c_0}{2} \right) \right) \\
&= c_0 + \left(\frac{1}{\kappa^{+}} \int_{\omega \in C_1} h^{\min} d\omega \right) \frac{c_1 - c_0}{2} \\
&= c_0 + \left(\frac{\gamma g^{\min}}{\kappa^{+}} \right) \frac{(c_1 - c_0)^2}{2} \\
&\geq c_0 + \left(\frac{\gamma g^{\min}}{\kappa^{+}} \right) \left(\frac{\kappa^{+}}{\gamma g^{\max}} \right)^2 \frac{1}{2} \\
&= c_0 + \frac{g^{\min} \kappa^{+}}{2\gamma (g^{\max})^2}.
\end{aligned}$$

Finally we find a lower bound on κ^{+}. Recall that

$$\kappa^{+} = \frac{\gamma (1 - G(p^*))}{k^{+}} = \frac{\gamma (1 - G(p^*))}{\max \left\{ 1, \left\lfloor \frac{\gamma(1 - G(p^*))}{\bar{\kappa}} \right\rfloor \right\}}.$$

If $1 > \left\lfloor \frac{\gamma(1-G(p^*))}{\bar{\kappa}} \right\rfloor$ then

$$\kappa^{+} = \gamma (1 - G(p^*)) \geq \gamma \left(1 - G\left(p^{NE}\right)\right),$$

provided that $p^* < p^{NE}$. If $1 < \left\lfloor \frac{\gamma(1-G(p^*))}{\bar{\kappa}} \right\rfloor$ then

$$\kappa^{+} = \frac{\gamma (1 - G(p^*))}{\left\lfloor \frac{\gamma(1-G(p^*))}{\bar{\kappa}} \right\rfloor} \geq \frac{\gamma (1 - G(p^*))}{\frac{\gamma(1-G(p^*))}{\bar{\kappa}}} = \bar{\kappa}.$$

Thus

$$\mathbb{E}\left[\omega|\omega \in C_1\right] \geq c_0 + \frac{g^{\min}}{2\gamma\left(g^{\max}\right)^2}\min\left\{\bar{\kappa}, \gamma\left(1 - G\left(p^{NE}\right)\right)\right\}.$$

■

A.3 Limit Case: $k^- \to \infty$ and k^+ Finite

Consider the case of a coarse categorization below p^*, as introduced in subsection 4.1.7. As before, buyers form a categorization using a heuristic that aims to split the mass of transactions into equally sized bins. The minimal probability mass in each category is $\bar{\kappa} \leq 1$. Due to the discontinuity of the probability density function of transactions h at p^*, the intervals $[0, p^*]$ and $(p^*, 1]$ are categorized separately. Below p^* the number of categories is

$$k^- := \max\left\{1, \left\lfloor \frac{1}{\bar{\kappa}} \int_0^{p^*} h(s)\, ds \right\rfloor\right\} = \max\left\{1, \left\lfloor \frac{G(p^*)}{\bar{\kappa}} \right\rfloor\right\},$$

each with a density of $\kappa^- = \frac{G(p^*)}{k^-}$. Above p^* the maximal number of categories is k^+ each with a density of κ^+, as defined above. Note that a categorization C is described completely by the three parameters k^-, k^+, and p^*. Thus, if we take the quality distribution G as given, the categorization C is completely determined by $\bar{\kappa}$, γ, and p^*.

We now show how the case of a perfect understanding of quality for ask prices less than or equal to p^* can be obtained by letting the number of categories below p^* go to infinity. Let $\bar{\kappa} \to 0$ and $\gamma \to 0$ at such a rate that $\frac{\bar{\kappa}}{\gamma} \to a$ for some constant $a > 0$. Note that, for any $p^* \in (0, 1)$,

$$k^+ \leq \max\left\{1, \frac{\gamma\left(1 - G\left(p^*\right)\right)}{\bar{\kappa}}\right\} \leq \max\left\{1, \frac{\gamma}{\bar{\kappa}}\right\}.$$

Thus $\frac{\bar{\kappa}}{\gamma} \to a$ implies that $1 \leq k^+ \leq 1/a$ for all[1] $p^* \in (0, 1)$. Intuitively, this corresponds to a situation with sufficiently many data on transactions to allow for an arbitrarily fine-grained categorization below p^* but a sufficiently low level of perfectly informed buyers

[1]More formally, consider sequences $\{\bar{\kappa}_n\}_{n=1}^{\infty}$ and $\{\gamma_n\}_{n=1}^{\infty}$ such that $\bar{\kappa}_n \to 0$ and $\gamma_n \to 0$, with $\frac{\bar{\kappa}_n}{\gamma_n} \to a$ for some constant $a > 0$. For each $n \in \mathbb{N}$, let $k_n^-(p^*)$ and $k_n^+(p^*)$ be induced by $\bar{\kappa}_n$ and γ_n, given p^*. For any n and p^*,

$$k_n^+(p^*) \leq \max\left\{1, \frac{\gamma_n\left(1 - G\left(p^*\right)\right)}{\bar{\kappa}_n}\right\} \leq \max\left\{1, \frac{\gamma_n}{\bar{\kappa}_n}\right\}.$$

It follows that for any K^- there is an N such that if $n > N$ then $k_n^+ \leq 1/a$ and $k_n^- > K^-$.

γ, and so a coarse categorization is required above[2] p^*. If $k^- \to \infty$ then for $p \leq p^*$,

$$\begin{aligned}\sum_{j=1}^{i(p)-1} \left(G(c_j) - G(c_{j-1})\right)\left(v\left(C_{i(p)}\right) - p\right) &\to \int_{\omega=0}^{p} (\omega + b - p)\, g(\omega)\, d\omega \\ &= \int_{\omega=0}^{p} \omega g(\omega)\, d\omega + (b-p)\, G(p) \\ &= G(p) \int_{\omega=0}^{p} \omega \frac{g(\omega)}{G(p)} d\omega + (b-p)\, G(p) \\ &= G(p)\left(\mathbb{E}\left[\omega|\omega \leq p\right] + b - p\right).\end{aligned}$$

Thus

$$\pi^{CE}(p|p^*) \to \begin{cases} G(p)\left(\mathbb{E}\left[\omega|\omega \leq p\right] + b - p\right) & \text{if } p \leq p^* \\ \\ G(p^*)\left(\mathbb{E}\left[\omega|\omega \leq p^*\right] + b - p\right) \\ + \sum_{j=1}^{i(p)-1} \left(G(c_j) - G(c_{j-1})\right)\left(v(C_j) - p\right) & \text{if } p > p^*. \\ + \left(G(p) - G\left(c_{i(p)-1}\right)\right)\left(v\left(C_{i(p)}\right) - p\right) \end{cases}$$

A.4 Centralized Market

In order to show that the dynamic $p_{t+1}^* = \mathbb{E}\left[\omega|\omega \leq p_t^*\right] + b$ converges to the steady-state p^C from any initial condition $p_0^* \in \left(\underline{p}, \overline{p}\right)$, it is sufficient so show that $\frac{\partial}{\partial p}\left(\mathbb{E}\left[\omega|\omega \leq p\right] + b - p\right) < 0$, or equivalently $\frac{\partial}{\partial p}\mathbb{E}\left[\omega|\omega \leq p\right] < 1$, for all $p \in \left(\underline{p}, \overline{p}\right)$. Since $\mathbb{E}\left[\omega|\omega \leq p^C\right] + b = p^C$, the condition $\frac{\partial}{\partial p}\mathbb{E}\left[\omega|\omega \leq p\right] < 1$ implies that if $p_t^* \in \left(\underline{p}, p^C\right)$, then $p_{t+1}^* = \mathbb{E}\left[\omega|\omega \leq p_t^*\right] + b > p_t^*$ and the condition $\frac{\partial}{\partial p}\mathbb{E}\left[\omega|\omega \leq p\right] > 0$ implies that $p_{t+1}^* < p^C$. Similarly, if $p_t^* \in \left(p^C, \overline{p}\right)$, then $p_{t+1}^* = \mathbb{E}\left[\omega|\omega \leq p_t^*\right] + b \in \left(p^C, p_t^*\right)$.

A.4.1 Local Stability

We have

$$\mathbb{E}\left[\omega|\omega \leq p\right] = \frac{1}{G(p)} \int_0^p \omega g(\omega)\, d\omega,$$

and so

$$\frac{\partial}{\partial p}\mathbb{E}\left[\omega|\omega \leq p\right] = -\frac{g(p)}{G(p)^2} \int_0^p \omega g(\omega)\, d\omega + \frac{1}{G(p)} p g(p) = \frac{g(p)}{G(p)}\left(p - \mathbb{E}\left[\omega|\omega \leq p\right]\right).$$

Note that

$$\left.\frac{\partial}{\partial p}\mathbb{E}\left[\omega|\omega \leq p\right] - 1\right|_{p=p^C} = \frac{g\left(p^C\right)}{G\left(p^C\right)}\left(p^C - \mathbb{E}\left[\omega|\omega \leq p^C\right]\right) - 1 = \frac{g\left(p^C\right)}{G\left(p^C\right)}(b) - 1.$$

[2] We may further restrict attention to $k^- \to \infty$ and $k^+ = 1$. For any $p^* \in (0,1)$ this holds when we let $a > 1$. In this case the categorization is arbitrarily fine-grained below p^* but maximally coarse above p^*.

This is negative if $\frac{g\left(p^C\right)}{G(p^C)} < \frac{1}{b}$, which is the case since (i) $p^C > p^{NE}$, and (ii) $\frac{g\left(p^{NE}\right)}{G(p^{NE})} = \frac{1}{b}$ and (iii) $\frac{\partial}{\partial p}\left(\frac{g(p)}{G(p)}\right) < 0$. By continuity, $\frac{\partial}{\partial p}\mathbb{E}\left[\omega|\omega \leq p\right] < 1$, for all p in an open interval around p^C.

A.4.2 Global Stability

As noted above, the log-concavity of G implies $\frac{\partial}{\partial p}\left(\frac{g(p)}{G(p)}\right) < 0$. Moreover, $(p - \mathbb{E}\left[\omega|\omega \leq p\right])$ is monotonically increasing in p if and only if G is log-concave, by Lemma 5 in Bagnoli and Bergstrom (2005). Note that, via integration by parts,

$$\begin{aligned} p - \mathbb{E}\left[\omega|\omega \leq p\right] &= p - \frac{1}{G(p)}\int_0^p \omega g(\omega)\,d\omega \\ &= p - \frac{1}{G(p)}\left(pG(p) - \int_0^p G(\omega)\,d\omega\right) \\ &= \frac{1}{G(p)}\int_0^p G(\omega)\,d\omega. \end{aligned}$$

Thus $\frac{\partial}{\partial p}\mathbb{E}\left[\omega|\omega \leq p\right] < 1$ is equivalent to (1).

SUPPLEMENT to Cycling and Categorical Learning in Decentralized Adverse Selection Economies

Philippe Jehiel and Erik Mohlin

S.1 Coarse Categorization below p^*

We pursue analytical results assuming a uniform quality distribution g. Hence each category $C_i = [c_{i-1}, c_i]$ below p^* will have the same width $c_i - c_{i-1} = p^*/k^-$ (where $k^- = \max\{1, \lfloor p^*/\bar{\kappa} \rfloor\}$) and each category $C_i = [c_{i-1}, c_i]$ above p^* will have the same width $c_i - c_{i-1} = \gamma(1-p^*)/k^+$ (where $k^+ = \max\{1, \lfloor \gamma(1-p^*)/\bar{\kappa} \rfloor\}$). The expected value for a buyer from trades with an ask price in C_i is $v(C_i) = (c_i + c_{i-1})/2 + b$, and the (perceived) expected utility of a buyer is

$$\pi^{CE}(p|p^*) = \sum_{j=1}^{i(p)-1} (c_j - c_{j-1})(v(C_j) - p) + \left(p - c_{i(p)-1}\right)\left(v\left(C_{i(p)}\right) - p\right).$$

S.1.1 Analytical Results

Convergence We can show that there is convergence to an interval containing the Nash price b. For this we need two lemmata.

Lemma S1 *There is some $\varepsilon^- > 0$ such that if $p^* < b - \frac{p^*}{2k^-}$ then $\frac{\partial \pi^{CE}(p)}{\partial p} > 0$ for all $p \leq p^* + \varepsilon$ (whenever $\frac{\partial \pi^{CE}(p)}{\partial p}$ is defined).*

Proof. For $p \in (c_{i-1}, c_i)$ with $c_i \leq c_{i(p^*)} = p^*$ we have

$$\begin{aligned}\frac{\partial \pi^{CE}(p)}{\partial p} &= \frac{c_i + c_{i-1}}{2} + b - 2p \\ &\geq \frac{c_i + c_{i-1}}{2} + b - 2c_i \\ &= b - c_i - \frac{c_i - c_{i-1}}{2} \\ &= b - c_i - \frac{p^*}{2k^-},\end{aligned}$$

which is positive since $p^* < b - \frac{p^*}{2k^-}$.

For $p \in \left(c_{i(p^*)}, c_{i(p^*)+1}\right) = \left(p^*, c_{i(p^*)+1}\right)$ we have

$$\lim_{p \downarrow p^*} \frac{\partial \pi^{CE}(p)}{\partial p} = \frac{c_{i(p^*)} + c_{i(p^*)+1}}{2} + b - 2p^*$$
$$= b - p^* + \frac{c_{i(p^*)+1} - c_{i(p^*)}}{2}$$
$$= b - p^* + \frac{(1 - p^*)}{2k^+},$$

which is positive since $p^* < b - \frac{p^*}{2k^-}$ implies that $p^* < b + \frac{(1-p^*)}{2k^+}$.

Recall that $\pi^{CE}(p|p^*)$ is concave within each category, and note that for $p \in \left(c_{i(p^*)}, c_{i(p^*)+1}\right)$

$$\frac{\partial \pi^{CE}(p)}{\partial p} = \frac{c_{i(p^*)} + c_{i(p^*)+1}}{2} + b - 2p = 0,$$

which implies that

$$p = \frac{1}{2}\left(\frac{c_{i(p^*)} + c_{i(p^*)+1}}{2} + b\right)$$
$$= \frac{1}{2}\left(p^* + \frac{c_{i(p^*)} - c_{i(p^*)+1}}{2} + b\right)$$
$$= \frac{1}{2}\left(p^* + \frac{(1 - p^*)}{2k^+} + b\right).$$

Thus, either $\frac{\partial \pi^{CE}(p)}{\partial p} > 0$ for all $p \in \left(c_{i(p^*)}, c_{i(p^*)+1}\right]$, or

$$\arg \max_{p \in \left(c_{i(p^*)}, c_{i(p^*)+1}\right]} \pi^{CE}(p|p^*) = \frac{1}{2}\left(p^* + \frac{(1 - p^*)}{2k^+} + b\right) \in \left(c_{i(p^*)}, c_{i(p^*)+1}\right].$$

Note that $p^* < b - \frac{p^*}{2k^-}$ implies

$$\frac{1}{2}\left(p^* + \frac{(1 - p^*)}{2k^+} + b\right) > \frac{1}{2}\left(p^* + \frac{(1 - p^*)}{2k^+} + p^* + \frac{p^*}{2k^-}\right)$$
$$= p^* + \frac{1}{2}\left(\frac{(1 - p^*)}{2k^+} + \frac{p^*}{2k^-}\right)$$
$$> p^* + \frac{1}{2}\left(\frac{(1 - p^*)\bar{\kappa}}{2} + \frac{\bar{\kappa} p^*}{2}\right)$$
$$= p^* + \frac{1}{4}\bar{\kappa},$$

where the last inequality uses

$$k^+ = \max\left\{1, \left\lfloor \frac{\gamma (1 - G(p^*))}{\bar{\kappa}} \right\rfloor\right\} \leq \max\left\{1, \frac{1}{\bar{\kappa}}\right\} = \frac{1}{\bar{\kappa}},$$

and

$$k^{-} = \max\left\{1, \left\lfloor \frac{G(p^*)}{\bar{\kappa}} \right\rfloor\right\} \leq \max\left\{1, \frac{1}{\bar{\kappa}}\right\} = \frac{1}{\bar{\kappa}}.$$

■

Lemma S2 *There is some* $\varepsilon^{+} > 0$ *such that if* $p^* > b + \frac{(1-p^*)}{2k^+}$ *then* $\frac{\partial \pi^{CE}(p)}{\partial p} < 0$ *for all* $p \geq p^* - \varepsilon$ *(whenever* $\frac{\partial \pi^{CE}(p)}{\partial p}$ *is defined).*

Proof. For $p \in (c_{i-1}, c_i)$ with $c_{i-1} \geq c_{i(p^*)} = p^*$ we have

$$\begin{aligned} \frac{\partial \pi^{CE}(p)}{\partial p} &= \frac{c_i + c_{i-1}}{2} + b - 2p \\ &\leq \frac{c_i + c_{i-1}}{2} + b - 2c_{i-1} \\ &= b - c_{i-1} + \frac{c_i - c_{i-1}}{2} \\ &\leq b - p^* + \frac{(1-p^*)}{2k^+}, \end{aligned}$$

which is positive since $p^* > b + \frac{(1-p^*)}{2k^+}$.

For $p \in \left(c_{i(p^*)-1}, c_{i(p^*)}\right) = \left(c_{i(p^*)-1}, p^*\right)$ we have

$$\begin{aligned} \lim_{p \uparrow p^*} \frac{\partial \pi^{CE}(p)}{\partial p} &= \frac{c_{i(p^*)-1} + c_{i(p^*)}}{2} + b - 2p^* \\ &= b - p^* - \frac{c_{i(p^*)} - c_{i(p^*)-1}}{2} \\ &= b - p^* - \frac{p^*}{2k^-}, \end{aligned}$$

which is positive since $p^* > b + \frac{(1-p^*)}{2k^+}$ implies that $p^* > b - \frac{p^*}{2k^-}$.

Recall that $\pi^{CE}(p|p^*)$ is concave within each category, and note that, for $p \in \left(c_{i(p^*)-1}, c_{i(p^*)}\right)$,

$$\frac{\partial \pi^{CE}(p)}{\partial p} = \frac{c_{i(p^*)-1} + c_{i(p^*)}}{2} + b - 2p = 0,$$

which implies that

$$\begin{aligned} p &= \frac{1}{2}\left(\frac{c_{i(p^*)-1} + c_{i(p^*)}}{2} + b\right) \\ &= \frac{1}{2}\left(p^* - \frac{c_{i(p^*)-1} - c_{i(p^*)}}{2} + b\right) \\ &= \frac{1}{2}\left(p^* - \frac{p^*}{2k^-} + b\right). \end{aligned}$$

Thus, either $\frac{\partial \pi^{CE}(p)}{\partial p} > 0$ for all $p \in \left(c_{i(p^*)-1}, c_{i(p^*)}\right]$, or

$$\arg \max_{p \in \left(c_{i(p^*)-1}, c_{i(p^*)}\right]} \pi^{CE}\left(p|p^*\right) = \frac{1}{2}\left(p^* - \frac{p^*}{2k^-} + b\right) \in \left(c_{i(p^*)-1}, c_{i(p^*)}\right].$$

Note that $p^* > b + \frac{(1-p^*)}{2k^+}$ implies that

$$\begin{aligned}
\frac{1}{2}\left(p^* + \frac{p^*}{2k^-} + b\right) &< \frac{1}{2}\left(p^* - \frac{p^*}{2k^-} + p^* - \frac{(1-p^*)}{2k^+}\right) \\
&= p^* - \frac{1}{2}\left(\frac{p^*}{2k^-} + \frac{(1-p^*)}{2k^+}\right) \\
&< p^* - \frac{1}{2}\left(\frac{\bar{\kappa} p^*}{2} + \frac{(1-p^*)\,\bar{\kappa}}{2}\right) \\
&= p^* - \frac{1}{4}\bar{\kappa},
\end{aligned}$$

where the last inequality uses $k^-, k^+ \leq \frac{1}{\bar{\kappa}}$, as argued at the end of the proof of Lemma S1. ■

We can now prove convergence to a an interval containing the Nash price b.

Proposition S1 *Suppose that $\bar{\kappa} < \frac{1}{2}$. Starting from any initial p_0^*, as $t \to \infty$ the dynamic will converge to, and never escape, the set*

$$P^*\left(\bar{\kappa}, \gamma\right) = \left(b - \frac{1}{2\left(\frac{1}{\bar{\kappa}} - 1\right)}, b + \frac{1}{2\left(\frac{\gamma}{\bar{\kappa}} - 1\right)}\right).$$

Proof. Note that

$$\begin{aligned}
k^- &= \max\left\{1, \left\lfloor \frac{p^*}{\bar{\kappa}} \right\rfloor\right\} \geq \max\left\{1, \frac{p^*}{\bar{\kappa}} - 1\right\}, \\
k^+ &= \max\left\{1, \left\lfloor \frac{\gamma\left(1-p^*\right)}{\bar{\kappa}} \right\rfloor\right\} \geq \max\left\{1, \frac{\gamma\left(1-p^*\right)}{\bar{\kappa}} - 1\right\}.
\end{aligned}$$

This further implies (using $\frac{1}{2} > \bar{\kappa}$ in the last two equalities to get rid of the min-operator),

$$\begin{aligned}
\frac{p^*}{2k^-} &\leq \min\left\{\frac{p^*}{2\left(\frac{p^*}{\bar{\kappa}}-1\right)}, \frac{p^*}{2}\right\} = \min\left\{\frac{1}{2\left(\frac{1}{\bar{\kappa}}-\frac{1}{p^*}\right)}, \frac{1}{2}\right\} \\
&\leq \min\left\{\frac{1}{2\left(\frac{1}{\bar{\kappa}}-1\right)}, \frac{1}{2}\right\} = \frac{1}{2\left(\frac{1}{\bar{\kappa}}-1\right)} \\
\frac{(1-p^*)}{2k^+} &\leq \min\left\{1, \frac{(1-p^*)}{2\left(\frac{\gamma}{\bar{\kappa}}(1-p^*)-1\right)}\right\} = \min\left\{1, \frac{1}{2\left(\frac{\gamma}{\bar{\kappa}}-\frac{1}{(1-p^*)}\right)}\right\} \\
&\leq \min\left\{1, \frac{1}{2\left(\frac{\gamma}{\bar{\kappa}}-1\right)}\right\} = \frac{1}{2\left(\frac{\gamma}{\bar{\kappa}}-1\right)}.
\end{aligned}$$

Thus, if $p^* < b - \frac{1}{2\left(\frac{1}{\bar{\kappa}}-1\right)}$ then $p^* < b - \frac{p^*}{2k^-}$ and hence S1 states that there is some $\varepsilon^- > 0$ such that $\frac{\partial \pi^{CE}(p)}{\partial p} > 0$ for all $p \leq p^* + \varepsilon^-$ (except at finitely many points of non-differentiability). Similarly, if $p^* > b + \frac{1}{2\left(\frac{\gamma}{\bar{\kappa}}-1\right)}$ then $p^* > b + \frac{(1-p^*)}{2k^+}$ and hence Lemma S2 states that there is some $\varepsilon^+ > 0$ such that $\frac{\partial \pi^{CE}(p)}{\partial p} < 0$ for all $p \geq p^* - \varepsilon^+$ (except at finitely many points of non-differentiability). This means that if $p^* < b - \frac{1}{2\left(\frac{1}{\bar{\kappa}}-1\right)}$ then $p^*_{t+1} - p^*_t > \varepsilon^-$ and if $p^* > b + \frac{1}{2\left(\frac{\gamma}{\bar{\kappa}}-1\right)}$ then $p^*_{t+1} - p^*_t < -\varepsilon^+$. ■

No Rest Point As the minimal mass of categories $\bar{\kappa}$ decreases and approaches 0 the set $P^*(\bar{\kappa}, \gamma)$ converges to (the singleton set consisting of) b. One might conjecture that $p^* = b$ is a steady-state for a small enough value of $\bar{\kappa}$. However, this is not the case.

Proposition S2 *Suppose that $p^* = b$. Then,*

$$\arg\max_{p\in[0,1]} \pi^{CE}(p) \leq b - \frac{b}{4k^-},$$

or

$$\arg\max_{p\in[0,1]} \pi^{CE}(p) \geq b + \frac{1-b}{4k^+}.$$

Proof. The constrained optimal $p \in \left(c_{i(p^*)-1}, c_{i(p^*)}\right]$ is given by

$$\begin{aligned}
\frac{\partial \pi^{CE}(p)}{\partial p} &= \frac{c_{i(p^*)} + c_{i(p^*)-1}}{2} + b - 2p = 0 \\
&\Longleftrightarrow \frac{b + b - \frac{p^*}{k^-}}{2} + b - 2p = 0 \\
&\Longleftrightarrow p = b - \frac{p^*}{4k^-},
\end{aligned}$$

and note that $b-\frac{p^*}{4k^-}\in\left(c_{i(p^*)-1},c_{i(p^*)}\right]$. The constrained optimal $p\in\left(c_{i(p^*)},c_{i(p^*)+1}\right]$ is given by

$$\begin{aligned}\frac{\partial\pi^{CE}(p)}{\partial p}&=\frac{c_{i(p^*)+1}+c_{i(p^*)}}{2}+b-2p=0\\ &\Longleftrightarrow\frac{b+\frac{1-p^*}{k^+}+b}{2}+b-2p=0\\ &\Longleftrightarrow p=b+\frac{1-p^*}{4k^+},\end{aligned}$$

and note that $b+\frac{1-p^*}{4k^+}\in\left(c_{i(p^*)},c_{i(p^*)+1}\right]$. ■

Cycle for $k^-=1$ and $k^+=1$ To obtain analytical results regarding cycles we now look at the special cases where $k^-=1$ and $k^+=1$. In these cases $v(C_1)=\frac{p^*}{2}+b$ and $v(C_2)=\frac{2-p^*}{2}+b$. We identify the following cycle with two price levels.

Proposition S3 *There is a price cycle with two prices, $p^{*low}=(2b+1/5)/3$ and $p^{*high}=(2b+4/5)/3$, such that, if $p_t^*=p^{*low}$ then $p_{t+1}^*=p^{*high}$, and if $p_t^*=p^{*high}$ then $p_{t+1}^*=p^{*low}$. Moreover, if $p_t^*>\bar{p}:=(2b+1/2)/3$ then $p_{t+1}^*<p_t^*$, whereas if $p_t^*<\bar{p}$ then $p_{t+1}^*>p_t^*$.*

Proof. We look for a cycle with two price levels. The first derivative at $p\in[0,p^*)$ is

$$\frac{\partial\pi^{CE}(p)}{\partial p}=\frac{p^*}{2}+b-2p,$$

and so the FOC for an interior constrained optimal $p\in[0,p^*]$ is $p=(p^*/2+b)/2$, generating utility $((p^*/2+b)/2)^2$.

The first derivative at $p\in(p^*,1]$ is

$$\frac{\partial\pi^{CE}(p)}{\partial p}=\frac{1+p^*}{2}+b-2p,$$

and so the FOC for an interior constrained optimal $p\in(p^*,1]$ is $p=((1+p^*)/2+b)/2$, generating utility $(((1+p^*)/2+b)/2)^2-p^*/2$. Thus the low category is perceived to contain the optimum if

$$\left(\frac{1}{2}\left(\frac{p^*}{2}+b\right)\right)^2-\left(\left(\frac{1}{2}\left(\frac{1+p^*}{2}+b\right)\right)^2-\frac{1}{2}p^*\right)=\frac{1}{16}(6p^*-4b-1)>0$$

or equivalently

$$p^*>\bar{p}:=\frac{2b+1/2}{3}.$$

This implies that if $p_t^*>\bar{p}$ then $p_{t+1}^*<p_t^*$, whereas if $p_t^*<\bar{p}$ then $p_{t+1}^*>p_t^*$.

In order to have a cycle with two price levels we need

$$p^{*low} = \frac{1}{2}\left(\frac{p^{*high}}{2} + b\right)$$

and

$$p^{*high} = \frac{1}{2}\left(\frac{1 + p^{*low}}{2} + b\right)$$

or equivalently

$$p^{*low} = \frac{1}{2}\left(\frac{1}{2}\left(\frac{1}{2}\left(\frac{1 + p^{*low}}{2} + b\right)\right) + b\right) \Longleftrightarrow p^{*low} = \frac{1 + 10b}{15} = \frac{2b + 1/5}{3}$$

and

$$p^{*high} = \frac{10b + 4}{15} = \frac{2b + 4/5}{3}.$$

■

If $b = 0.3$ then $p^{*high} = 0.467$ and $p^{*low} = 0.267$. This is illustrated in Figure S1 ($\gamma = 0.1$, $\bar{\kappa} = 1/2$, $b = 0.3$), for initial prices equal to 0.1, 0.5, and 0.9, respectively.

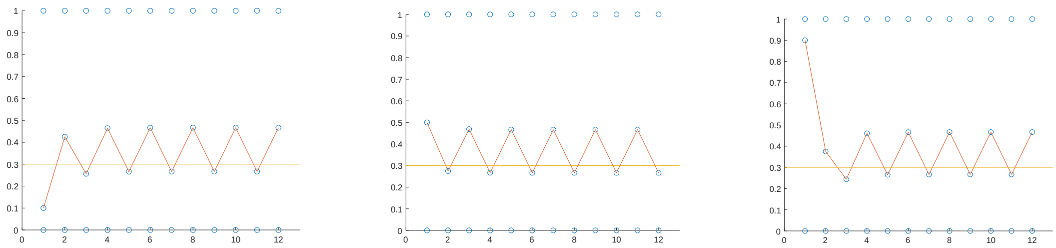

$p_0 = 0.1$ $p_0 = 0.5$ $p_0 = 0.9$

Figure S1

S.1.2 Numerical Results

We expect our main result, regarding convergence to a price cycle, to be relevant when there are sufficiently many categories below p^* and sufficiently few categories above p^*. We examine this numerically by examining how price movements vary when we vary the information parameter γ, and the categorization coarseness parameter $\bar{\kappa}$. Throughout we let $b = 0.3$ and use an initial price of $p = 0.1$. The results are displayed in Figure S2. The qualitative pattern does not depend on the exact level of b, and convergence does not depend on the initial price. In general, increasing the share of informed buyers γ and decreasing the required mass per category $\bar{\kappa}$, dampens fluctuations by increasing the number of categories above p^*.

$\bar{\kappa} = 1/10$ $\bar{\kappa} = 1/30$ $\bar{\kappa} = 1/50$

$\gamma = 0.05$

$\gamma = 0.1$

$\gamma = 0.2$

$\gamma = 0.3$

Figure S2

S.1.3 Categorization without (Automatic) Category Boundary at p^*

Now that categories are used both below and above p^* one may wonder what happens if buyers ignore the discontinuity at p^* when forming categories. Suppose that each category has a mass of transactions exactly equal to $\bar{\kappa}$, regardless of whether the category is below p^*, above p^*, or includes p^*. That is,

$$c_i - c_{i-1} \begin{cases} = \bar{\kappa} & \text{if } c_i \leq p^* \\ \in \left(\bar{\kappa}, \frac{\bar{\kappa}}{\gamma}\right) & p^* \in (c_{i-1}, c_i] \\ = \frac{\bar{\kappa}}{\gamma} & \text{if } c_{i-1} \geq p^* \end{cases} .$$

For the case of a uniform g we can show that starting from any initial p_0^*, as $t \to \infty$ the dynamic will converge to, and never escape, the set

$$\tilde{P}^* (\bar{\kappa}, \gamma) = \left(b - \frac{2\bar{\kappa}}{\gamma}, b + \frac{2\bar{\kappa}}{\gamma}\right).$$

As the probability of perfect information γ increases, and as the number of categories k increases, the set $\tilde{P}^* (\bar{\kappa}, \gamma)$ converges to (the singleton set consisting of) b. Still, as before, we can show that b is not a rest point.

Figure S3 reports numerical results for the dynamics in the case where there is no (automatic) category boundary at p^*. We vary γ and $\bar{\kappa}$ in the same way as in the numerical investigation of the case with a category boundary at p^*. In some cases there is pronounced cycling, while in other cases there appears to be none. The reason is that cycles appear for parameter combinations that generate a category boundary sufficiently close to p^* (for each p^* in the cycle). In the case of $\bar{\kappa} = 1/10$ with $\gamma = 0.05$ or 0.1 the trajectory depends on the initial condition. For consistency, we report the trajectory starting at $p = 0.1$.

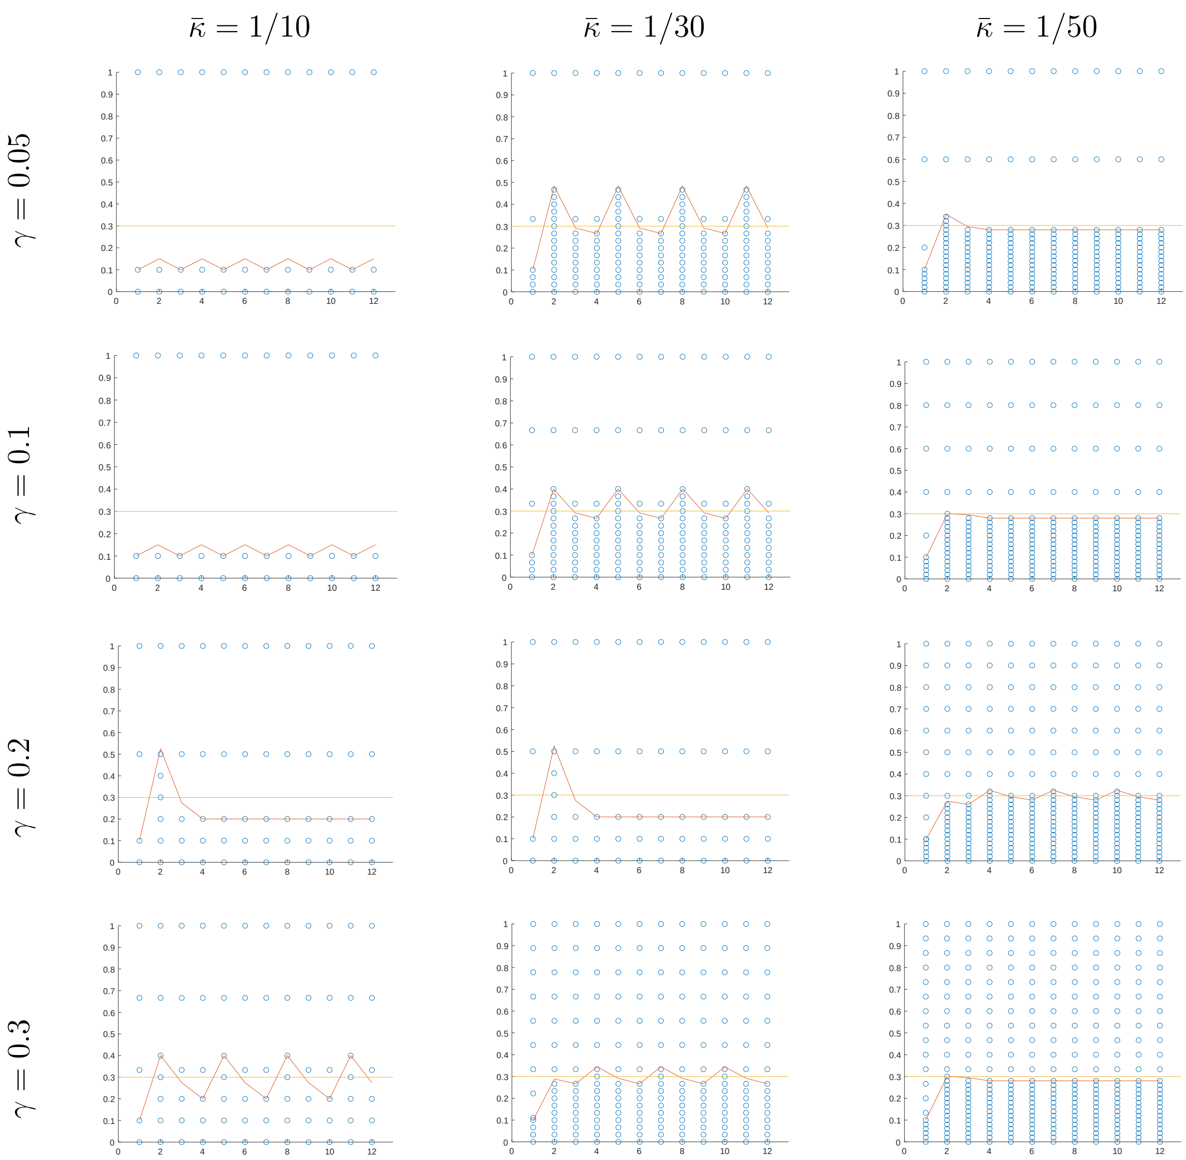
$\bar{\kappa} = 1/10$
$\bar{\kappa} = 1/30$
$\bar{\kappa} = 1/50$
$\gamma = 0.05$
$\gamma = 0.1$
$\gamma = 0.2$
$\gamma = 0.3$

Figure S3

S.2 Multiplicative Gains from Trade

Note that, as in the additive case, we still have

$$\max_{p \in C_i} \pi^{CE}(p) = \max_{p \in C_i} \sum_{j=1}^{i(p)-1} \left(G(c_j) - G(c_{j-1})\right)\left(v(C_j) - p\right) + \left(G(p) - G\left(c_{i(p)-1}\right)\right)\left(v\left(C_{i(p)}\right) - p\right),$$

and the first derivative at $p \in \left(c_{i(p)-1}, c_{i(p)}\right)$ is

$$\frac{\partial \pi^{CE}(p)}{\partial p} = g(p)\left(\mathbb{E}\left[\omega | c_{i(p)-1} < \omega \leq c_{i(p)}\right] \cdot \beta - p\right) - G(p).$$

S.2.1 Nash Equilibrium

Proposition S4 *The unique Nash equilibrium is* $p^{NE} = 0$.

Proof. In the Nash equilibrium the buyers have correct expectations about the mapping between ask price and quality. They maximize

$$\pi^{NE}(p) = \int_{\omega=0}^{p} (\omega \cdot \beta - p)\, g(\omega)\, d\omega = G(p)\left(\mathbb{E}[\omega | \omega \leq p]\, \beta - p\right).$$

From the additive case we know that

$$\frac{\partial}{\partial p}\left(\mathbb{E}[\omega | \omega \leq p]\right) = \frac{g(p)}{G(p)}\left(p - \mathbb{E}[\omega | \omega \leq p]\right).$$

Thus

$$\begin{aligned}
&\frac{\partial}{\partial p}\left(G(p)\left(\mathbb{E}[\omega|\omega \leq p] \cdot \beta - p\right)\right) \\
&= g(p)\left(\mathbb{E}[\omega|\omega \leq p] \cdot \beta - p\right) + G(p)\left(\frac{\partial}{\partial p}\left(\mathbb{E}[\omega|\omega \leq p]\right) \cdot \beta - 1\right) \\
&= g(p)\left(\mathbb{E}[\omega|\omega \leq p] \cdot \beta - p\right) + G(p)\left(\frac{g(p)}{G(p)}\left(p - \mathbb{E}[\omega|\omega \leq p]\right) \cdot \beta - 1\right) \\
&= g(p)\left(\mathbb{E}[\omega|\omega \leq p] \cdot \beta - p\right) + g(p)\left(p - \mathbb{E}[\omega|\omega \leq p] \cdot \beta\right) - G(p) \\
&= g(p)\left(\mathbb{E}[\omega|\omega \leq p] \cdot \beta - p\right) - g(p)\left(\mathbb{E}[\omega|\omega \leq p] \cdot \beta - p\right) - G(p) \\
&= -G(p) < 0.
\end{aligned}$$

■

S.2.2 Nash Not a Rest Point

Lemma S3 *If* $p_t^* = p^{NE} = 0$ *then,*

$$p_{t+1}^* = \arg\max_{p \in [0,1]} \pi^{CE}(p | p_t^*) > 0.$$

Proof. Differentiating π^{CE} at $p \in (c_0, c_1)$, and letting p go to $0 = c_0$, we obtain

$$\left.\frac{\partial \pi^{CE}(p)}{\partial p}\right|_{p \downarrow p^{NE}=0} = g(p)\left(v\left(C_{i(p)}\right) - p\right) - G(p)\Big|_{p=p^{NE}=0}$$
$$= g(0)\left(\mathbb{E}\left[\omega | c_{i-1} < \omega \leq c_i\right]\beta\right) > 0.$$

■

S.2.3 Cycle When $k^- \to \infty$ and $k^+ < \infty$

Define

$$p^{\beta(NE)} := \arg\max_{p \in [0,1]} \pi^{CE}(p | p^* = 0).$$

Let $k^- = \frac{p^{\beta(NE)}}{\bar{\kappa}} \to \infty$ and $k^+ = \frac{\gamma\left(1-p^{\beta(NE)}\right)}{\bar{\kappa}} \to a < \infty$. This means that for any $p^* \geq p^{\beta(NE)}$ we have

$$k^- = \frac{p^*}{\bar{\kappa}} > \frac{p^{\beta(NE)}}{\bar{\kappa}} \to \infty$$

and

$$k^+ = \frac{\gamma(1-p^*)}{\bar{\kappa}} = \frac{\gamma\left(1-p^{\beta(NE)}\right)}{\bar{\kappa}} \frac{(1-p^*)}{\left(1-p^{\beta(NE)}\right)} \to a \frac{(1-p^*)}{\left(1-p^{\beta(NE)}\right)} < \infty.$$

This implies that, for any $p^* \geq p^{\beta(NE)}$ and any $p \leq p^*$, we have

$$\sum_{j=1}^{p} (c_j - c_{j-1})\left(v\left(C_{i(p)}\right) - p\right) \to \int_{\omega=0}^{p} (\omega\beta - p)\, g(\omega)\, d\omega = G(p)\left(\mathbb{E}[\omega | \omega \leq p]\beta - p\right),$$

and so, for any $p^* \geq p^{\beta(NE)}$

$$\pi^{CE}(p|p^*) \to \begin{cases} G(p)\left(\mathbb{E}[\omega|\omega \leq p]\beta - p\right) & \text{if } p \leq p^*, \\ \\ G(p^*)\left(\mathbb{E}[\omega|\omega \leq p^*]\beta - p\right) & \\ +\sum_{j=i(p^*)+1}^{i(p)-1}\left(G(c_j) - G(c_{j-1})\right)\left(v(C_j) - p\right) & \text{if } p > p^*. \\ +\left(G(p) - G\left(c_{i(p)-1}\right)\right)\left(v\left(C_{i(p)}\right) - p\right) & \end{cases}$$

Suppose that at time t it was the case that $p_t^* = p^{NE} = 0$. Then at time $t+1$ the buyers bid $p_{t+1}^* = p^{\beta(NE)} > p$. For period $t+2$, the following result demonstrates that either $p_{t+2}^* = p^{NE} = 0$ or $p_{t+2}^* > p_{t+1}^*$.

Lemma S4 *If $p_t^* > p^{NE} = 0$, then either*

$$p_{t+1}^* = \arg\max_{p \in [0,1]} \pi^{CE}(p) = p^{NE} = 0$$

or

$$p^*_{t+1} = \arg\max_{p\in[0,1]} \pi^{CE}(p) > p^*_t.$$

Proof. Since $\pi^{CE}(p)$ coincides with $\pi^{NE}(p)$ on $\left[0, p^*_{t+1}\right]$, the constrained optimal $p \in \left[0, p^*_{t+1}\right]$ is at $p = p^{NE} = 0$. Suppose that $\arg\max_{p\in\left[p^*_{t+1},1\right]} \pi^{CE}(p) = p^*_{t+1}$ (requiring $\left.\frac{\partial \pi^{CE}(p)}{\partial p}\right|_{p\downarrow p^*_{t+1}} \leq 0$). Then by continuity of $\pi^{NE}(p)$, $\arg\max_{p\in[0,1]} \pi^{CE}(p) = p^{NE} = 0$. ■

We can now prove that there is a cycle.

Proposition S5 *There exists an increasing sequence* $(p^{(1)}, ..., p^{(m)})$ *with* $m \geq 2$ *and* $p^{(1)} = p^{NE}$ *such that if* $p^*_t = p^{(i)}$ *for* $i \in \{1, ..., m-1\}$ *then* $p^*_{t+1} = p^{(i+1)}$*, and if* $p^*_t = p^{(m)}$ *then* $p^*_{t+1} = p^{(1)}$*. Moreover, the dynamic converges to the set* $\left\{(p^{(1)}, ..., p^{(m)})\right\}$ *from any initial price* p_0*.*

Proof. Assume that there is no cycle. Lemma S4 implies that

$$p^*_{t+1} = \arg\max_{p\in[0,1]} \pi^{CE}(p|p^*_t) > p^*_t,$$

for all t. Since $p^*_t \leq 1$ for all t, it follows that $p^*_t \to \bar{p}$ for some $\bar{p} > p^{NE}$ as $t \to \infty$. This implies that $\left|p^*_{t+1} - p^*_t\right| \to 0$, which, by continuity of $\pi^{CE}(p|p^*_t)$, implies that

$$\left|\pi^{CE}\left(p^*_{t+1}|p^*_t\right) - \pi^{CE}\left(p^*_t|p^*_t\right)\right| \to 0.$$

Since $\pi^{CE}(p|p^*_t) = \pi^{NE}(p)$ for $p \in [0, p^*_t]$, we have $\left|\pi^{CE}\left(p^*_{t+1}|p^*_t\right) - \pi^{NE}\left(p^*_t\right)\right| \to 0$, and consequently $\pi^{CE}\left(p^*_{t+1}|p^*_t\right) \to \pi^{NE}(\bar{p})$. Since the Nash equilibrium p^{NE} is unique it holds that $\pi^{NE}\left(p^{NE}\right) > \pi^{NE}(\bar{p})$, and since $\pi^{CE}(p|p^*_t) = \pi^{NE}(p)$ for $p \in [0, p^*_t]$ we get

$$\pi^{CE}\left(p^*_{t+1}|p^*_t\right) \to \pi^{NE}(\bar{p}) < \pi^{NE}\left(p^{NE}\right) = \pi^{CE}\left(p^{NE}|p^*_t\right).$$

This is in contradiction to $p^*_{t+1} = \arg\max_{p\in[0,1]} \pi^{CE}(p|p^*_t)$. We conclude that there is a cycle. Lemma S3 and Lemma S4 imply that the cycle consists of p^{NE} and one or more prices above p^{NE}.

Note that the preceding argument can be used to show, that there is convergence to the cycle, from which there is no escape. To see this suppose (to obtain a contradiction) that there is some $p^*_1 > p^{NE}$ that does not belong to the cycle (i.e., $p^*_1 \neq p^{(1)}$ for all $i \in \{1, ..., m\}$), from which there is no convergence to the cycle. This means that $p^*_{t+1} > p^*_t$ for all t and $p^*_t \to \bar{p}$ for some $\bar{p} \in \left[p^{NE}, p^{(m)}\right]$ as $t \to \infty$. ■

S.3 Buyer Competition

A rest point of the dynamic with Bertrand pricing, denoted by p^B, must satisfy $\pi_i^B\left(p^B, p^B|p^B\right) = 0$, i.e., satisfy $\pi_i^{CE}\left(p^B|p^B\right) = 0$ or, equivalently,

$$G\left(p^B\right)\left(\mathbb{E}\left[\omega|\omega \leq p^B\right] + b - p^B\right) = 0 \Longleftrightarrow p^B = \mathbb{E}\left[\omega|\omega \leq p^B\right] + b. \tag{S1}$$

In order for $p_t^* = p^B$ to be a rest point it must be the case that $\pi_i^B\left(p_i, p^B|p_t^B\right) \leq 0$ for all $p_i > p^B$ (since otherwise buyer competition would lead to increasing prices). For the same reasons as in the case without buyer competition, $\pi_i^B\left(p_i, p^B|p_t^B\right)$ is continuous in p_i and concave in p_i within each category. Thus it holds that $\pi_i^B\left(p_i, p^B|p_t^B\right) \leq 0$ for all $p_i > p^B$ if and only if $\lim_{p_i \downarrow p^B} \frac{\partial}{\partial p_i}\pi_i^B\left(p_i, p^B|p_t^B\right) \leq 0$.

For simplicity, suppose that there is a single category above p_t^* (i.e., $\bar{\kappa}$ sufficiently large and γ sufficiently small). For $p_i > p^B$ we have

$$\frac{\partial}{\partial p_i}\pi_i^B\left(p_i, p^B|p_t^B\right) = g\left(p_i\right)\left(\mathbb{E}\left[\omega|p^B \leq \omega \leq 1\right] + b - p_i\right) - G\left(p_i\right),$$

and so

$$\lim_{p_i \downarrow p^B}\frac{\partial}{\partial p_i}\pi_i^B\left(p_i, p^B|p_t^B\right) = g\left(p^B\right)\left(\mathbb{E}\left[\omega|p^B \leq \omega \leq 1\right] + b - p^B\right) - G\left(p^B\right).$$

Using (S1) this is non-positive iff

$$\frac{g\left(p^B\right)}{G\left(p^B\right)}\left(\mathbb{E}\left[\omega|p^B \leq \omega \leq 1\right] - \mathbb{E}\left[\omega|\omega \leq p^B\right]\right) \leq 1. \tag{S2}$$

If instead $\lim_{p_i \downarrow p^B} \frac{\partial}{\partial p_i}\pi_i^B\left(p_i, p^B|p_t^B\right) > 0$ then p^B is not a rest point. From $p_t^* = p^B$ the dynamic will move to p_{t+1}^*, which solves $\pi_i^B\left(p_{t+1}^*, p_{t+1}^*|p_t^*\right) = 0$, i.e., p_{t+1}^* that solves

$$G\left(p_t^*\right)\left(\mathbb{E}\left[\omega|\omega \leq p_t^*\right] + b - p_{t+1}^*\right) + \left(G\left(p_{t+1}^*\right) - G\left(p_t^*\right)\right)\left(\mathbb{E}\left[\omega|p_t^* \leq \omega \leq 1\right] + b - p_{t+1}^*\right) = 0. \tag{S3}$$

For uniform g equation (S2) becomes $p^B \geq 1/4$. Moreover, in the uniform case (S1) becomes $p^B = 2b$. Thus, in order for the perceived marginal utility of p_i to be non-positive at p^B we need $2b \geq 1/4 \Longleftrightarrow b \geq 1/8$. If instead $b < 1/8$ then when $p_{t+1}^* = 2b$ equation (S3) becomes

$$\left(p_{t+1}^*\right)^2 - \frac{1}{2}p_{t+1}^* - 4b^2 + b = 0.$$

The solutions are $p_{t+1}^* = 2b$ and $p_{t+1}^* = \frac{1}{2} - 2b$. With the former solution $p_{t+1}^* = p_t^*$. With the latter solution (which requires $b < \frac{1}{4}$ to be non-negative) $p_{t+1}^* - p_t^* = \frac{1}{2} - 4b$. This is positive only if $b < 1/8$. Next note that if $p_{t+1}^* > p_t^*$ then $\pi_i^B\left(p_{t+1}^*, p_{t+1}^*|p_{t+1}^*\right) < 0$. Thus we must have $p_{t+2}^* > p_{t+1}^*$ or $p_{t+1}^* = p_t^* = p^B$. (If $p_t^* = \frac{1}{2} - 2b$ then $\lim_{p_i \downarrow p_t^*} \frac{\partial}{\partial p_i}\pi_i^B\left(p_i, p_t^*|p_t^*\right) =$

$4b - \frac{1}{4}$, which is positive iff $b > 1/16$.)

Summing up, we have shown, for the uniform case, that if $b < 1/8$, and if there is always a single category above p^* (i.e., $\bar{\kappa}$ sufficiently large and γ sufficiently small) then there is a cycle, such that if $p_t^* = p^B = 2b$ then $p_{t+1}^* > \frac{1}{2} - 2b$. There may be more than two prices in the cycle (if $b > \frac{1}{16}$) but eventually the dynamic returns to p^B.

S.4 Uniform g

Proposition S6 *Suppose that g is uniform. If $p_t^* = p^{NE} = b$ then*

$$p_{t+1}^* = \frac{1+3b}{4} = \arg\max_{p\in[0,1]} \pi^{CE}(p)$$

and

$$p_{t+2}^* = b = \arg\max_{p\in[0,1]} \pi^{CE}(p).$$

Proof of Proposition S6. If $k^+ = 1$ then

$$\pi^{CE}(p) \to \begin{cases} p\left(\frac{p}{2} + b - p\right) & \text{if } p \leq p^* \\ \begin{array}{c} p^*\left(\frac{p^*}{2} + b - p\right) \\ +\left(p - c_{i(p^*)}\right)\left(\frac{1+p^*}{2} - p\right) \end{array} & \text{if } p > p^*. \end{cases}$$

Suppose that $p_t^* = p^{NE} = b = c_{i(p^*)}$. The first derivative at $p \in [0, p^*)$ is $\frac{\partial \pi^{CE}(p)}{\partial p} = b - p$, and so, since $p^* = b$, the constrained optimal $p \in [0, p^*] = [0, b]$ is $p^* = b$. The constrained optimal $p \in \left(c_{i(p^*)}, 1\right]$ is given by

$$\frac{\partial \pi^{CE}(p)}{\partial p} = \frac{1+p^*}{2} + b - 2p = 0 \Longleftrightarrow p = \frac{1+3b}{4}.$$

Note that $\lim_{p\uparrow b} \frac{\partial}{\partial p}\pi^{CE}(p) = 0$ and

$$\lim_{p\downarrow b} \frac{\partial \pi^{CE}(p)}{\partial p} = \lim_{p\downarrow b}\left(2b + \frac{1+b}{2} - 2p\right) = \frac{1+b}{2} > 0.$$

Thus, since $\pi^{CE}(p)$ is continuous, we have

$$\max_{p\in\left(c_{i(p^*)},1\right]} \pi^{CE}(p) = \pi^{CE}\left(\tilde{p}_{C_{i(p^*)+1}}\right) > \max_{p\in[0,p^*]} \pi^{CE}(p).$$

Now suppose that $p_{t+1}^* = (1+3b)/4 = c_{i(p^*)}$. The first derivative at $p \in [0, p^*)$ is $\frac{\partial}{\partial p}\pi^{CE}(p) = b - p$, and so, since $p^* > b$, the constrained optimal $p \in [0, p^*]$ is b, yielding a payoff of $b^2/2$. For $p \in (p^*, 1]$ we have

$$\frac{\partial \pi^{CE}(p)}{\partial p} \geq \lim_{p\downarrow p^*} \frac{\partial \pi^{CE}(p)}{\partial p} = \frac{1+p^*}{2} + b - 2p^* = \frac{1}{8}(1-b) > 0$$

and the optimum is given by the FOC

$$\frac{\partial \pi^{CE}(p)}{\partial p} = \frac{1+p^*}{2} + b - 2p = 0 \Longleftrightarrow p = \frac{11b+5}{16},$$

which yields the payoff

$$\pi^{CE}\left(\frac{11b+5}{16}\right)=\frac{1}{256}\left(137b^2-2b-7\right).$$

We compare this with the constrained optimal $p\in[0,p^*]$, which is b, yielding a payoff of $b^2/2$. Since

$$\begin{aligned}\frac{b^2}{2}-\frac{1}{256}\left(137b^2-2b-7\right)&=\frac{1}{256}\left(2b+7-9b^2\right)>\frac{1}{256}\left(2b+7-9b\right)\\&=\frac{7}{256}\left(1-b\right)>0,\end{aligned}$$

we conclude that if $p^*=\left(1+3b\right)/4=c_{i(p^*)}$ then

$$\max_{p\in[0,p^*]}\pi^{CE}\left(b\right)>\max_{p\in(p^*,1]}\pi^{CE}\left(\frac{11b+5}{16}\right).$$

■

S.5 Comparison with Related Models

In this section we compare our model with a number of related models. Throughout we assume a uniform quality distribution g.

S.5.1 Cursed Equilibrium

Eyster and Rabin (2005) define a notion of cursedness, capturing an inability to (fully) understand the mapping between types and actions in an incomplete information game. A *fully cursed* buyer fails entirely to take the selection into account and believes that the expected value is $\mathbb{E}[v]$ independent of the ask price, and so the perceived expected payoff is

$$\begin{aligned}\pi^{Cursed}(p) &= \Pr(\omega \leq p) \cdot (\mathbb{E}[\omega + b] - p) \\ &= p\left(\frac{1}{2} + b - p\right).\end{aligned}$$

The solution is $p^{Cursed} = \frac{1}{2}\left(b + \frac{1}{2}\right)$. Thus a cursed buyer bids too much if $b < 1/2$ and too little if $b > 1/2$. The reason for the former effect is that a cursed buyer is overoptimistic about the value v since $\mathbb{E}[v] \geq E[v|\omega \leq p]$ for all p. The reason for the latter effect is that the cursed buyer fails to understand that a higher bid will elicit higher quality.

A *partially cursed* buyer partially takes selection into account. Her expected payoff is

$$\begin{aligned}\pi^{Partial}(p) &= (1-\delta)\,\pi^{NE}(p) + \delta\pi^{Cursed}(p) \\ &= p\left((1-\delta)\left(\frac{p}{2}\right) + \delta\left(\frac{1}{2}\right) + b - p\right).\end{aligned}$$

The solution is $p = \frac{1}{(1+\delta)}\left(b + \frac{\delta}{2}\right)$, which for $\delta = 1$ reduces to, $\frac{1}{2}\left(b + \frac{1}{2}\right)$, the choice of a fully cursed buyer, and for $\delta = 0$ reduces to b, the Nash equilibrium choice.

S.5.2 Behavioral Equilibrium

Esponda (2008) defines a notion of behavioural equilibrium. Buyers have unlimited information about all past bid and ask prices as well as information about whether trade occurred. Thus buyers are able to perfectly assess the probability of trade as a function of bid price. Buyers have information about quality in all those cases trade took place, but no information about quality in the cases trade did not occur. Thus in equilibrium buyers are able to form correct beliefs about how quality depends on ask prices below the equilibrium price, but not on those above the equilibrium price. In order to form a belief about quality associated with bid prices above the equilibrium price they have to use data on bid prices below the equilibrium price. Esponda considers a naive and a sophisticated way of doing this.

Let p^N be the equilibrium bid price with *naive* buyers. It is w.l.o.g. to assume that $p^N \leq 1$. Naive buyers believe that the expected value is the same above the equilibrium price as it is below the equilibrium price. Formally, the buyers assume that $\mathbb{E}\left[v|a \leq p^N\right] = \mathbb{E}\left[v|a > p^N\right]$, and hence their expected payoff is

$$\pi^N(p) = \Pr(a \leq p) \cdot \left(\mathbb{E}\left[\omega + b|a \leq p^N\right] - p\right) = p \cdot \left(\frac{p^N}{2} + b - p\right).$$

This is concave in p and so the FOC gives us

$$\max_p \pi^N(p) = \frac{1}{2}\left(\frac{p^N}{2} + b\right).$$

In equilibrium we must have $p^N = \max_p \pi^N\left(|p^N\right)$. The solution is $p^N = \frac{2}{3}b < p^{NE}$, so the adverse-selection problem is aggravated.

Esponda also considers *sophisticated* buyers. Let p^S be the equilibrium price in this case. It is w.l.o.g. to assume that $p^S \leq 1$. Let $\rho^S(p)$ denote the expected quality conditional on a bid price p being accepted. Since pairs consisting of an ask prices and a quality are observed for ask prices below the equilibrium bid price p^S, is it is assumed that $\rho^S(p) = \mathbb{E}[\omega + b|\omega \leq p]$ for $p \leq p^S$. Moreover, $\rho^S(p)$ is assumed to be non-decreasing in p for all p^S. That is, sophisticated buyers may have some understanding that higher bids elicit higher quality. This implies that for $p > p^S$ we have $\rho^S(p) \geq \mathbb{E}\left[\omega + b|\omega \leq p^S\right]$. (In addition Esponda adds restrictions to the effect that beliefs are consistent with observed distributions of quality.) The expected payoff is given by

$$\begin{aligned}
\pi^S(p) &= \Pr(\omega \leq p) \cdot \left(\rho^S(p) - p\right) \\
&= \begin{cases} \Pr(\omega \leq p)\left(\mathbb{E}[\omega + b|\omega \leq p] - p\right) & p \leq p^S \\ \Pr(\omega \leq p)\left(\rho^S(p) - p\right) & p > p^S \end{cases} \\
&= \begin{cases} p\left(b - \frac{p}{2}\right) & p \leq p^S \\ p\left(\rho^S(p) - p\right) & p > p^S \end{cases}.
\end{aligned}$$

Esponda notes that p^N is a lower bound for the set of equilibrium prices with a sophisticated buyer. He also notes that p^{NE} is an upper bound for the set of equilibrium prices with a sophisticated buyer. To see this, note that if $p^S > p^{NE}$ then the buyer knows $\mathbb{E}[v|\omega \leq p]$ for all $p < p^S$ and so will realize that he is better off at a price below p^S. Alternatively, the buyer will want to deviate to an even higher price, which means that p^S is not an equilibrium.

S.5.3 Analogy-based expectations equilibrium (ABEE)

Suppose that, as in Esponda's setup, buyers only observe quality when trade took place, and suppose that, as in our model above, with probability γ a buyer is perfectly informed about quality. In analogy-based expectations equilibrium (ABEE), due to Jehiel (2005) and Jehiel and Koessler (2008), buyers use categories to judge quality as a function of prices (as in our model) but categories are fixed. For simplicity, assume that all categories have the same width $1/k$. We may either assume (i) that players observe quality only in cases where trade took place and that with probability γ a buyer is perfectly informed about quality, or (ii) that players observe quality also when trade did not take place. In either case the payoff function has the same form as in our model:

$$\pi^{ABEE}(p) = \sum_{j=1}^{i(p)-1} (c_j - c_{j-1})(v(C_j) - p) + (p - c_{i(p)-1})(v(C_{i(p)}) - p).$$

One can show that as the number of categories goes to infinity the solution to the problem of maximizing $\pi^{ABEE}(p)$ on $[0,1]$ approaches the Nash equilibrium price $b = p^{NE}$. If there is a single category then the solution coincides with the fully cursed equilibrium price $p^{Cursed} = \frac{1}{2}\left(b + \frac{1}{2}\right)$.

Proposition S7 *As the number of categories goes to infinity the solution to* $\max_{p\in[0,1]} \pi^{ABEE}(p)$ *approaches the Nash equilibrium price* $b = p^{NE}$. *If there is a single category then the solution to* $\max_{p\in[0,1]} \pi^{ABEE}(p)$ *coincides with the fully cursed equilibrium price* $p^{Cursed} = \frac{1}{2}\left(b + \frac{1}{2}\right)$.

Proof. Note that, for $p \in (c_{i-1}, c_i)$,

$$\begin{aligned}\frac{\partial \pi^{CE}(p)}{\partial p} &= \frac{c_i + c_{i-1}}{2} + b - 2p \\ &\geq \frac{c_i + c_{i-1}}{2} + b - 2c_i \\ &= b - c_i - \frac{c_i - c_{i-1}}{2} \\ &= b - c_i - \frac{1}{2k}.\end{aligned}$$

Thus, if $b - \frac{1}{2k} \geq c_i$ then $\frac{\partial \pi^{CE}(p)}{\partial p} > 0$ for all $p \in (c_{i-1}, c_i)$. Furthermore, note that, that for $p \in (c_{i-1}, c_i)$,

$$\begin{aligned}\frac{\partial \pi^{CE}(p)}{\partial p} &\leq \frac{c_i + c_{i-1}}{2} + b - 2c_{i-1} \\ &= b - c_{i-1} + \frac{c_i - c_{i-1}}{2} \\ &= b - c_{i-1} + \frac{1}{2k}.\end{aligned}$$

Thus, if $b + \frac{1}{2k} \leq c_{i-1}$ then $\frac{\partial \pi^{CE}(p)}{\partial p} > 0$ for all $p \in (c_{i-1}, c_i)$. It follows that as the number of categories goes to infinity the solution to $\max_{p \in [0,1]} \pi^{ABEE}(p)$ approaches the Nash price $b = p^{NE}$.

By contrast, if there is a single category then we obtain the same payoff function as for a fully cursed buyer, namely,

$$\pi^{ABEE}(p) = p\left(\frac{1}{2} + b - p\right) = \pi^{Cursed}(p),$$

and so the solution is $p^{ABEE} = p^{Cursed} = \frac{1}{2}\left(b + \frac{1}{2}\right)$. ■