The Effect of Potential Collusion on Equilibrium Prices*

Tom-Reiel Heggedal[†], Espen R. Moen[‡], and Magnus V. Knutsen[§]

February 20, 2024

Abstract

We investigate how potential collusion, interpreted as a positive prior probability that collusion may occur, influences equilibrium prices in states of the world in which collusion, in fact, does not occur. We explore the mechanism in a theoretical model of consumer search, based on the Stahl (1989) framework, and show that with potential collusion, equilibrium prices are, in fact, higher even in the absence of collusion. We explore the mechanism in a series of laboratory experiments. The results are qualitatively consistent with our theoretical predictions.

*The research is financed by BI Norwegian Business School.
[†]BI Norwegian Business School & CESAR, : tom-reiel.heggedal@bi.no
[‡]BI Norwegian Business School & CESAR & CEPR: espen.r.moen@bi.no
[§]BI Norwegian Business School & CESAR: magnus.v.knutsen@bi.no

1 Introduction

Collusion, by definition, leads to higher prices. The focus in this paper is on how potential collusion, interpreted as a positive prior probability that collusion occurs, may influence equilibrium prices in states of the world in which collusion in fact does not occur. Fear of collusion may influence the search behaviour of consumers. In particular, potential collusion may reduce consumers' incentives to continue search after being offered a high price. Non-colluding sellers may take advantage of this, and set higher prices than when there is no potential for collusion.

We first explore this mechanism in a theoretical model of consumer search based on the Stahl (1989) framework. The model contains two crucial features for our mechanism to work. First, buyers believe that collusion may take place with a strictly positive probability. Second, buyers are not fully informed about the prices offered by all sellers and may search to obtain more price quotes.

In the simplest version of our model, there is an exogenous probability that sellers may collide on a price p^M. This price is higher than any price offered in equilibrium in the absence of potential collusion and would not have been accepted by any consumers. In the presence of potential collusion, the situation is different. When a consumer visits a seller that posts p^M, she updates the probability that firms collude and buys the good with a strictly positive probability. This in turn induces some sellers to set p^M even in the absence of collusion, and the entire price distribution shifts upwards. As a result, equilibrium prices are higher even when collusion, in fact, does not occur.

In the second part of this paper, we explore the mechanism in a laboratory experiment. In our experiment, participants have roles as sellers or buyers. The probability that sellers collude varies between the treatments. When collusion occurs, prices are set exogenously at p^M for all sellers. We run a total of four treatments. In our control treatment, the exogenous probability of collusion is zero. In the three remaining treatments, the exogenous probability of collusion is 10%, 20% and 30%, respectively.

We find that potential collusion changes the behaviour of both sellers and buyers in a way that is qualitatively consistent with the predictions of the model. Regarding posted prices, we find that potential collusion leads to a mass point in the price distribution on the collusive price p^M and an upward shift in the remaining price distribution below p^M. Both of these findings are in line with theory. As a result, there is an increase in the average price compared to our control that ranges between 17 percent and 101 percent. Regarding search behavior, we find that search rates are low across all treatments, which is in accordance with our theoretical predictions.

Although there are deviations from equilibrium predictions in the data, both buyers and sellers best-respond remarkably well. Across our four treatments, the share of buyer decisions that are consistent with a best response, given the empirical reservation price, ranges between 91 and 97 percent. Furthermore, we find that the majority of posted prices are within the ranges that yield the highest profits. Thus, the pricing behaviour of sellers is to a large extent profit-maximismizing.

The major discrepancy from theory is that there are differences between treatments that are not in line with theory. In particular, there are numerous instances of price postings below the theoretical minimum in the low-probability treatment, and there is excessive mass on the collusive price in the high-probability treatment. We explore deviations from equilibrium predictions employing a model of reinforcement learning. We find that subjects are sensitive to the experience gained throughout the experiment and that this can account for differences from equilibrium predictions as well as differences between treatments.

1.1 Related Literature

Our paper contributes both to theoretical and experimental literatures on markets with search and markets with collusion. First, we contribute to the theoretical literature on consumer search with learning (of market state). In this literature firms typically observe costs privately and costs are positively correlated. In Benabou and Gertner (1993) identical consumers observe the price of one firm and make inference about the other firm's costs before making their search decision. In a similar environment Dana Jr (1994) includes informed and uninformed buyers, where the latter first observes one price for free and can then pay a cost to get fully informed about all prices.[1] Janssen et al. (2011) take this type of production cost uncertainty to a sequential search environment, similar to ours. They develop a Stahl (1989) type model with stochastic, perfectly correlated costs observed only by the firms, and show that both informed and uninformed consumers pay higher expected prices under cost uncertainty. Further properties of the model using different equilibrium refinements are analyzed in Janssen et al. (2017). In contrast, our paper focuses directly on collusion in prices and examines how the fear of such collusion may lead to increased prices in the absence of actual collusion. Furthermore, since our paper involves buyers making direct inferences on prices, there is no need for refinements on out-of-equilibrium beliefs to solve the model.

There is also a string of related search papers that includes consumers' inference about downstream firms' costs in vertical market contexts (see e.g., Janssen and Shelegia (2015), Lubensky (2017), Janssen (2020) and Janssen and Shelegia (2020)).[2]

Second, our papers contributes to the theoretical literature on consumer search and tacit collusion in repeated games. Nilsson et al. (1999) considers a repeated game version of Burdett and Judd (1983) and shows that lower search costs facilitate collusion as profits on the punishment path is lower. Campbell et al. (2005) corroborate this result in a Stahl (1989) framework, and show that the scope for collusion is larger if firms can monitor each other's prices. Petrikaitė (2016) study search costs and cartel stability in repeated game versions of Wolinsky (1986) and Stahl (1989). The main

[1] There are also some paper that takes this framework to a dynamic setting, analyzing how costs correlated over time affects pricing dynamics (see e.g., Yang and Ye (2008) and Tappata (2009)).

[2] There also are several papers on search where consumers learn about other states of the market than production costs. In Lauermann et al. (2012) traders gradually learn about market demand and supply through a sequence of multilateral bargaining rounds. Atayev (2022) studies uncertainty about product availability. Mauring (2017) and Mauring (2020) analyze how consumers learn about the price offer distribution from past searchers' trade decisions in an environment with non strategic sellers.

insight is that search costs may have two countervailing effects on stability. While higher search costs softens the punishment from deviating from the cartel, higher costs may also decrease (short term) deviation incentives pending on the details of the competitive environment. Shadarevian (2022) extends this framework to include private information in the firm side about positively correlated costs in the spirit of Benabou and Gertner (1993). Montag and Winter (2020) nests the models of Petrikaitė (2016) and Schultz (2017) to study how price transparency affects the stability of collusive behaviour. Obradovits and Plaickner (2020) considers a price-directed search model in which consumers observes prices but search for match quality.[3]

A common characteristic for these repeated games is that buyers are aware of whether prices are determined through play on the cooperative path or the punishment path. In contrast, in our paper buyers form beliefs about collusion, and these beliefs may diminish consumers' incentives to search. Exploiting this, non-colluding sellers set higher prices, even in the absence of actual collusion.

Third, our paper contributes to the experimental literature on posted-offers models with price dispersion. Cason et al. (2003) and Cason et al. (2021) examine the noisy sequential-search model of Burdett and Judd (1983); Cason and Datta (2006) and Cason and Mago (2010) examine the sequential search model of Robert and Stahl (1993) with advertising; Helland et al. (2017) examine the capacity constraints model of Lester (2011); Morgan et al. (2006) examine the simultaneous pricing model of Varian (1980) while and Heggedal et al. (2023) study sequential pricing in this framework.[4]

The zest of this literature is that the comparative statics is often in line with equilibrium predictions. However, point predictions can deviate substantially from equilibrium, in particular in experiments with human buyers (see Cason and Mago (2010)). Moreover, on the individual level, subjects typically do not mix over prices,

[3] Schultz (2005) and Schultz (2017) study tacit collusion in a Hotelling framework without search in which only a fraction of consumers observes both firms' prices akin to Varian (1980).

[4] There also exists a literature that investigates different aspects of the Diamond (1971) paradox in the lab. See for instance Grether et al. (1988), Davis and Holt (1996) and Abrams et al. (2000).

rather prices are correlated over time (see Cason et al. (2003) and Cason et al. (2021)). In our paper, we produce qualitative predictions, as well as several point predictions, in the lab using human subjects both as sellers and buyers in complex search markets.

Last, our paper contributes to the literature on collusion in the lab. To the best of our knowledge there are only two experimental studies of collusion in posted-offers models with price dispersion. Moellers et al. (2016) studies collusion in a repeated game variant of the Stahl (1989) model. They show that lower search costs do not lower average prices although buyers search more often, and that (cheap talk) communication between sellers increases prices and reduces search. Orzen (2008) studies tacit collusion a repeated game variant of the Varian (1980) model. He finds that with repeated interaction, duopolists post substantially higher prices than in the one-shot game, whereas prices in quadropolies remain very similar.

In contrast to these two papers on repeated games, we find that the potential for collusion in a one-shot-game leads to higher prices in the lab even in the absence of collusion.

There is a large experimental literature on collusion in repeated price-setting games without price-information frictions. This literature shows that the specifics of the market structure are crucial for whether collusion is successful in generating profits. See Potters and Suetens (2013) and Harrington Jr et al. (2016) for excellent surveys.[5]

The remainder of the paper is organized as follows: Section 2 outlines the theoretical framework, and section 3 presents our experimental design and procedures. In section 4 we report the results from our experiment, while section 5 discusses deviations from equilibrium predictions. Section 6 offers a brief conclusion.

[5] There is also a small literature on search and collusion in the field. In the study conducted by Moraga-González et al. (2023), an exploration of the car dealership market was undertaken using random search models. Although conclusive evidence of collusion was not discovered in the analyzed markets, the findings revealed that reducing consumers' search costs could result in elevated prices. Similarly, Nishida and Remer (2018) identified adverse effects associated with diminished search costs when investigating the retail gasoline market.

2 Mechanism and Model

Before we dive into the model, we will discuss the underlying economic mechanism, which we believe applies more broadly than in our rather narrow model environment as long as prices are influenced by consumer search.

Our mechanism is related to the beliefs and incentives of searching consumers. A core element is that (a subset of) consumers have limited information about the prices set by the different sellers, and that they can gather information about prices through costly search. If a consumer visits a seller and finds that it sets a high price, she will continue searching and visit another store if the expected gain from doing so is higher than the search cost. The consumers' search behaviour disciplines the sellers when setting prices and leads to lower equilibrium prices.

The risk of collusion may partly corrupt this mechanism. Suppose first that the consumers know that there is a probability that the sellers collude and set higher prices. A buyer who visits a store that charges a high price will not know if this high price reflects a high-price strategy of the seller, i.e., a seller-specific high price, or if it is due to collusion. If it is a seller-specific high-price strategy, the consumer may want to search again. However, if it reflects collusion, the gain from search is lower or non-existence, as the other sellers will charge a high price as well. Therefore, an *a priori* positive probability of collusion will reduce the incentives of buyers to continue searching after observing a high price.

This induced change in consumer search will in turn influence seller behaviour. If the sellers know that buyers search less due to the fear of collusion, they may be tempted to set a high price even in the absence of collusion, as the probability of being rejected by the consumers is lower. They may use the possibility of collusion as a "disguise" for a seller-specific high price. Hence, the possibility of collusion will induce some sellers to set a high price without collusion.

A common feature of price models with search frictions is that they generate price dispersion, and this gives rise to a source of amplification. As some sellers set a higher price to mimic price collusion, the competition among sellers setting lower prices soften. Hence, the entire price distribution shifts up, leading to even higher average prices.

In the rest of this section, we will explore these mechanisms in the so-called Stahl model, a work-horse model within consumer search.

Model set-up

We consider the Stahl model, with consumers' willingness to pay equal to 1, and with parameters such that the supremum of the price support in the Stahl equilibrium, denoted p_s, is strictly below 1. With exogenous probability x firms collude on a price $p^M > p_s$. Let ρ denote the probability that firms set p^M when there is no exogenous collusion. Let q denote the probability that a consumer searches if observing p^M. We focus on symmetric equilibria. Let $[p_0, p_1]$ denote the support of the continuous part of the distribution. Let c denote search costs of consumers if searching to obtain a second price quote. Let u denote the number (measure) of uninformed customers per seller and I the number (measure) of informed customers. The number of sellers is 2. The costs to the sellers are normalized to zero. We will derive the equilibrium of the model when collusion does not occur (when collusion occurs, sellers trivially set $p = p^M$).

We will distinguish between two equilibrium candidates. In the equilibrium candidate without consumer search, uninformed consumers never search after observing the monopoly price p^M. In the equilibrium candidate with consumer search, uninformed consumers search after observing p^M with an endogenous probability $q \in (0, 1]$. As will be clear below, which of the equilibrium candidates that constitute an equilibrium (ie, whether or not uninformed consumers will search or not after observing p^M), depends on the prior probability x of collusion.

Equilibrium candidate without consumer search

Suppose first that x is high so that consumers never search if they observe $p = p^M$ (below we show formally that consumers do not search for high values of x) . Let $F(p)$ be the probability distribution of p on $[p_0, p_1]$ contingent on $p \leq p_1$. Then the profit of a firm that sets a price at or below p_1 is

$$\pi(\rho) = p\left[u + \rho I + (1-\rho)I(1 - F(p))\right] = p\left[\tilde{u} + \tilde{I}(1 - F(p))\right] \tag{1}$$

where $\tilde{u} = u + \rho I$ and $\tilde{I} = (1-\rho)I$. Hence $F(p)$ is identical to the distribution with $u = \tilde{u}$ and $I = \tilde{I}$.

Equilibrium candidate without consumer search is a value $\rho \in [0, 1]$ and a distribution $F(p)$ with support $[p_0, p_1]$ satisfying the following conditions:

1. Equal profit when setting p^M and p_1:

$$p^M(u + \rho\frac{I}{2}) = p_1(u + \rho I) \tag{2}$$

2. Equal profits for all $p \in [p_0, p_1]$ implying that $F(p)$ given by

$$1 - F(p) = \frac{\tilde{u}(p_1 - p)}{p\tilde{I}} \tag{3}$$

   where p_0 is defined by $F(p_0) = 0$.

3. Consumers are indifferent between searching or not searching at p_1:

$$(1-\rho)\int_{p_0}^{p_1} F(p)dp = c \tag{4}$$

If a consumer observes p_1, she knows with certainty that collusion does not take place. Still, there is a probability ρ that the other firm sets p^M. Hence, for the buyer to be indifferent between searching and not searching, (13) must be satisfied.

Let us first make some observations. First, the lower part of the equilibrium is identical to the equilibrium of the Stahl model with $\bar{u}$ uninformed consumers per seller, $\bar{I}$ informed consumers, and search costs $\tilde{c}/(1-\rho)$. In particular, it follows that if $\rho > 0$, then $p_1 > p_s$.

Second, if (in contrast to our assumption) $p_s > p^M$, there cannot be a mass point at p^M. In this case, the standard argument against mass points applies; a firm setting p^M will be strictly better off by undercutting p^M slightly and attract all the informed customers if the opponent sets p^M (which happens with strictly positive probability).

Second, if $p^M > p_s$, then $p^M > p_1$ in equilibrium. To see this, note that if $\rho > 0$ and $p^M \leq p_1$, i.e., in an interval in which workers do not search, the undercutting argument from Varian (1980) again applies, and we cannot be in equilibrium. If $\rho = 0$, consumers will always know whether collusion takes place or not, and we must have $p_1 = p_s$.

Finally, we want to know how p_1 depends on ρ. For a given p_1, an increase in ρ decreases $\int_{p_0}^{p_1} F(p)dp$, so p_1 defined by (13) increases in ρ. Hence, there is a unique $\bar{\rho}$ such that (13) has a solution $p_1 \leq p^M$ if and only if $\rho \leq \bar{\rho}$.

Result 1. *Suppose $p_s < p^M$. Then the equilibrium candidate without consumer search exists and is unique.*

For any given $\bar{\rho}$, equation (12) and (13) uniquely define $F(p)$ and its support $[p_0, p_1]$ with $p_1 < p^M$. Define $\pi^1(\rho)$ as the profit at $p_1(\rho)$ and π^M the profit at the collusion price. The latter is independent of ρ. It is enough to show that $\pi^1(\rho)$ is strictly increasing in ρ, that $\pi^1(0) < \pi^M$, and that $\pi^1(\bar{\rho}) > \pi^M$. If $\rho = 0$, the distribution defined by (12) and (13) is as in the Stahl model, with a lower profit than at p^M. If ρ goes to $\bar{\rho}$, π^1 goes to $p^M(u+I) > p^M(u+I/2) = \pi^M$. That π^1 increases in ρ follows easily.

It should be noted that the equilibrium price distribution is independent of x, the probability that collusion takes place. When firms set prices, they *know* that collusion does not take place. When consumers observe a price different from p^M, they also know that collusion does not occur. Only if consumers observe p^M, they are uncertain whether collusion has taken place or not. In this section x is assumed to be so high that consumers do not search when observing p^M, and hence their beliefs do not influence equilibrium. In the next subsection, we relax this assumption.

Consumer search

A consumer observing p_1 is indifferent between searching and not. A consumer who observes p^M does not know if there is collusion or not, and this reduces the incentives to search. Given that a consumer observes a price p^M, the conditional probability that collusion takes place is given by $x^M = \frac{x}{x+(1-x)\rho}$ (by Bayes law), and the complementary probability (that collusion does not take place) is $1 - x^M = \frac{(1-x)\rho}{x+(1-x)\rho}$. If there is no collusion, the gain from search is $p^M - p^1$ plus the gain from search at p^1, which is c. Hence we must have that

$$(1 - x^M)\left(c + p^M - p_1\right) \leq c \tag{5}$$

or that

$$\frac{(1-x)\rho}{x + (1-x)\rho} \leq \frac{c}{c + p^M - p_1} \tag{6}$$

The left-hand side is strictly decreasing in x for a given ρ, and is 1 for $x = 0$ and 0 for $x = 1$. It is also strictly increasing in ρ, going from 0 to $1 - x$ as ρ goes from 0 to 1. Recall that the equilibrium values of ρ and p_1 are independent of x when there is no consumer search. Hence, for any equilibrium values (ρ^*, p_1^*) there exists a value $\bar{x}$ such that (6) is satisfied whenever $x \geq \bar{x}$

Result 2. *Consider an equilibrium candidate without consumer search, and let ρ^* and p_1^* denote the corresponding values of ρ and p_1. Then there exists a value $\bar{x} \in (0,1)$ such that consumers will not search at p^M whenever $x \geq \bar{x}$, where $\bar{x}$ is given by (6) with $\rho = \rho^*$ and $p_1 = p_1^*$.*

Equilibrium candidate with consumer search

Suppose then that (6) is not satisfied and that consumers observing p^M choose to search with probability q. We then have one new equation that must be satisfied, (6), and one more variable to be determined, q. Furthermore, q will influence the profitability of setting a low price, and hence the other equilibrium equations. Specifically, there will be more uniformed consumers searching. However, the *structure* of the equilibrium is the same as before. The price distribution has a mass point at p^M. Conditional on not setting p^M, the sellers set prices according to a continuous distribution function $F(p)$ with support $[p_0, p_1]$.

Consider a seller who sets a price $p \in [p_0, p_1]$. She will get u uninformed buyers directly. In addition, there is a probability ρ that the other seller sets p^M, attracting u uninformed buyers, and q of these (ρuq in expectation) will search again and end up with our seller. Finally, she will get all informed customers if and only if the other seller sets a price above p. Hence, the profit of this seller is

$$\pi(p) = p\left[(1+q\rho)u + \rho I + (1-\rho)I(1-F(p))\right] = p(\tilde{u} + \tilde{I}) \tag{7}$$

with $\tilde{u} = (1+q\rho)u + \rho I$ and $\tilde{I} = (1-\rho)I$ (as before), analogous with (1).

The *equilibrium candidate with consumer search* is defined as two values $\rho \in [0,1]$ and $q \in (0,1]$, and a continuous distribution function $F(p)$ on $[p_0, p_1]$, satisfying

1. Equal profit when setting p^M and p_1:

$$p^M(u(1-q) + I\rho/2) = p_1(u(1+q\rho) + \rho I) \tag{8}$$

2. Equal profits for all $p \in [p_0, p_1]$ implying that $F(p)$ given by

$$1 - F(p) = \frac{\tilde{u}(p_1 - p)}{p\tilde{I}} \tag{9}$$

with $\tilde{u} = (1+q\rho)u + \rho I$ and $\tilde{I} = (1-\rho)I$. p_0 is defined by $F(p_0) = 0$. This is the same equations as (12) above, but note that the definition of $\tilde{u}$ is different.

3. Consumers are indifferent between searching at not searching at p_1:

$$(1-\rho)\int_{p_0}^{p_1} F(p)dp = c \tag{10}$$

Again the equation is the same as the equilibrium equation (13), but the distribution function F is different.

4. Consumers at p^M are indifferent between searching and not searching: (6) is satisfied with equality.

A remark on the nature of the equilibrium is here in place. The equilibrium is sequential, in that sellers first set prices and then buyers respond, and both buyers and sellers play with mixed strategies at p^M. When visiting a seller that sets p^M, the strategy of the other seller makes the buyer indifferent between searching or not. The mixing of the competitor and the buyers makes a seller indifferent between the different prices. One may wonder why a seller does not slightly undercut p^M to ensure that buyers buy with certainty. The problem is that this strategy does not work, since the consumer then learns that there is no collusion and will continue to search with probability 1.

Result 3. *Suppose $x \leq \bar{x}$. Then the candidate equilibrium with consumer search exists and is unique and constitutes an equilibrium of the game.*

Sketch of proof: For a given ρ and q, it follows easily that $F(p)$ exists and is unique. Therefore, we can construct a mapping $\Gamma : [0,1] \times [0,1] \mapsto [0,1] \times [0,1]$ as follows: For a given pair (ρ, q) in the domain, calculate $p_1(\rho, q)$, which is continuous in both arguments. Given p_1, set ρ' by (6), if it exceeds $1-x$ set $\rho' = 1-x$. Then use (8) to calculate q', with 0 and 1 as bounds. This defines $(\rho', q') = \Gamma(\rho, q)$. The mapping is defined on the closed and compact set and is continuous; hence Brower's fixed-point theorem ensures that it has a fixed point. This fixed point, and the corresponding distribution $F(p)$ on $[p_0, p_1]$ is an equilibrium by construction.

Extensions

In the appendix we skectch an extension of the model, in which sellers, if collusion, set a common price according to a continuous price distribution G on an interval I^M above the maximum Stahl price. We show that without collusion, the sellers will randomise and set prices in I^M. The equilibrium price distribution contains no mass points. Furthermore, if consumers draw a price on the lower part of I, they randomise between searching and not searching. The equilibrium analysis is not yet complete.

3 Experimental design and procedures

The objective of our design is to empirically examine the core mechanism of our model. To achieve this goal, we compare pricing and searching decisions in two environments: one where consumers anticipate that the price may be set exogenously at the collusive level, and another where pricing is always endogenous.

We have four main treatments. In our baseline treatment, $T0$, the price is never set exogenously. In treatments $T10$, $T20$ and $T30$ the probability that the price is set exogenously is 0.10, 020 and 0.30, respectively. We implement parameterized versions of the model presented in section 2 in which there is no search in equilibrium to accentuate the difference in price setting between when $x = 0$ and $x > 0$. The random variable is independently and identically distributed across games.

3.1 Implementation

Subjects play 40 identical and independent games within a treatment. In each game, four subjects are matched into a market. Two of the subjects play the role as sellers, and two play the roles as the uninformed buyer. Sellers have zero marginal cost, while buyers have unit demand with a valuation of 100 experimental currency units (ECU). Roles remain fixed throughout the experiment.

The only treatment variation is the probability the price is set exogenously. At the beginning of each game, a probability draw determines whether sellers simultaneously and independently choose a price in the interval [0,100] ECU or the price is set exogenously at 70 ECU for all sellers.[6] Both sellers know the realization of the draw while buyers do not know this.

A buyer observes the price of one of the sellers in his/her market, while the other buyer observes the price of the other seller. Buyers can choose to purchase at the observed price or to pay a cost of 20 ECU to also observe the price of the other seller in the market.[7] This search decision is the only decision of buyers, purchasing is automated.[8] If a buyer pays to observe the price of the other seller, the buyer purchases from the other seller if and only if that seller has a lower price.

The concept of informed buyers is operationalized by automated additional sales. The seller with the lowest price in a market will make two such additional sales at her/his posted price. In case of price-ties, the sellers make one additional sale each. The additional sales do not depend on the decisions of buyers nor do they affect buyers' payoffs.

After each game, sellers and buyers receive feedback on prices set in their market and ECUs earned. Buyers are not informed about whether the price is set exogenously or by the sellers. After a game is completed, subjects are randomly re-matched into a new market. Matching is done within matching blocks of 8. Subjects are never exposed to more than one treatment (between-subject design). In our analysis we

[6] We picked 70 as the collusive price to both create sufficient distance from the upper bound p_1 and to be as far away from 1 to minimize impact from any fairness preferences on decisions.

[7] The search cost is set to give flexibility for the other parameters, in particular the probability of exogenous pricing. Note that the valuation (and price range) of 100 is not binding for our chosen parameters, and we could increase this number arbitrarily without changing any equilibrium predictions. See e.g. Hong and Shum (2006) for estimation of search cost using field data.

[8] We automate purchases to zoom in on the key decision of buyers and to avoid ordering effects in the sampling of prices.

regard average behavior within blocks as independent observations.

3.2 Hypotheses

Our main treatment measures are sellers' posted prices and buyers' purchase prices in games where the price is not set exogenously. We also measure buyers' search decisions. We use the model of section 3 to derive predictions for our four treatments. Given our chosen parameters, the no-search condition in equation 6 holds in all our treatments. Consequently, the model predicts no search in all treatments. As a result, theoretical price distributions, and therefore expected prices, are the same across treatments $T10$, $T20$ and $T30$, see figure 1 and table 1.

Figure 1: Theoretical price distributions by treatment

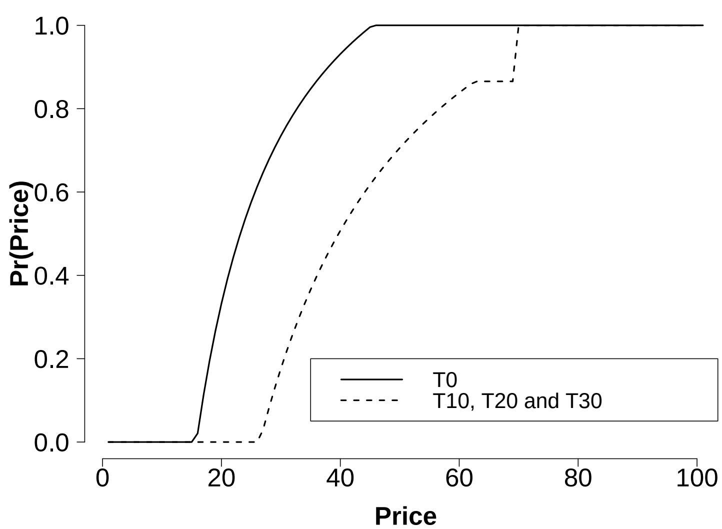

Table 1: Theoretical predictions

	$T0$	$T10$, $T20$ and $T30$
$E(p)$	24.4	42.6
p_0	14.8	26.5
p_1	44.4	62.6
ρ	0	0.13

Based on theoretical predictions, our primary hypothesis is that endogenous prices will be higher in treatments where prices may happen to be exogenous than in the baseline treatment, where this possibility is excluded. Our secondary hypothesis is that buyers will not search (or search little) across all treatments.

As the theoretical predictions are the same for treatments $T10$, $T20$ and $T30$, one may wonder why introducing different positive probabilities for the exogenous collusive price? There is a large literature showing that subjects in probabilistic environments often make different choices than predicted by standard Bayesian updating and expected utility maximization. In particular, updating is found to depend on priors in ways not consistent with by Baye's law.[9] Thus, there are reasons to believe that search decisions–and consequently rational price posting–may depend on the exogenous probability differently than what follows from bayesian updating in equilibrium of the model. After presenting the main results, we will explore deviations from the theoretical predictions in section **??**.

A pre-study plan, including a pilot study, for the experiment was posted at the AEA RCT registry on May 17th 2023.[10] The pre-study plan specifies treatments $T0$, $T10$ and $T20$, and we follow the plan in the set up of hypotheses and significance testing. $T30$ was added on to the study at a later stage.

The pilot study was carried out with two matching blocks for $T0$ and three matching blocks for $T20$. The average posted price of sellers was 25 ECU in $T0$ and 45

[9]See for instance Tversky and Kahneman (1971), Kahneman and Tversky (1972), Ouwersloot et al. (1998), Charness and Levin (2005), and Alós-Ferrer and Garagnani (2023)

[10]See https://www.socialscienceregistry. org/trials/11401.

ECU in $T20$ whereas the variances (between blocks) was 80 and 50, respectively. Only observations from the last 20 games was used.

Based on the average posted prices and the variance, we calculate the sample size needed to reach a power of 95 percent or better, given a 5 percent significance level and a Wilcoxon rank sum test. This estimate was obtained using the method described in Bellemare et al. (2016). The power-threshold required is reached with 6 independent matching blocks per treatment. Data from the pilot is included in the analysis of this paper.

3.3 Data collection

Data were collected in the Research Lab at BI Norwegian Business School in Oslo in the period April 2023 to June 2023. Subjects were recruited from the general student population of the BI Norwegian Business School.[11] Recruitment and subject management were administered through ORSEE (Greiner, 2015). On arrival, subjects were randomly allocated to cubicles (to break up social ties). Written instructions were handed out and read aloud by the administrator (to achieve public knowledge of the rules). A full set of instructions is provided in the supplementary online materials. All decisions were taken anonymously in a network of computers.

At the conclusion of a session, subjects were privately payed based on total ECU earned in two randomly drawn games. The exchange rate was 1 ECU = 2 NOK in all treatments. The protocol was implemented in zTree (Fischbacher, 2007). In total 192 participants participated in experiment sessions lasting on average one hour and earning on average 347 NOK.

[11] We argue that behavior in our sample is representative for decision makers in market contexts. Mounting evidence shows that behaviors in convenience samples (CSs) are generally representative of the general population; of students who do not self-select into lab experiments; and of workers in online labor markets, such as, Mechanical Turk (see Snowberg and Yariv (2021)). Moreover, behaviors in CSs often compare well to that of professionals, such as traders and managers (see Fréchette (2016) and Ball and Cech (1996)). Taking this research into account, we believe our results are informative of a mechanism that may also be important in real posted price markets.

4 Results

All results that are presented here exclude data from the 20 first games of the experiment. Also, we exclude prices from games in which the price was set exogenously.[12]

Table 2 provides a preview of our main result. In the treatments where the price is set exogenously with a strictly positive probability ($x > 0$), prices posted by subjects are significantly higher. In particular, posted prices in treatments $T10$, $T20$ and $T30$ are on average more than 50 percent higher than posted prices in treatment $T0$.

Table 2: Regression

	Posted Price
constant	27.62*** (2.85)
$I\{x > 0\}$	14.85*** (4.19)
N	1960

Note: First 20 games and games were the price was set exogenously are excluded. Standard errors (in parentheses) are clustered at the matching group level.

We divide the rest of our results in to four subsections: Section 4.1 presents the empirical price distributions, average prices and a comparison with equilibrium prices. Section 4.2 present results on search behavior and a comparison with equilibrium search behavior. In section 4.3 we determine the optimal search behavior given the empirical price distribution. In section D we calculate the optimal pricing strategy.

4.1 Prices

Figure 2 presents the empirical distribution of posted prices.[13] Consistent with equilibrium predictions, the price distributions in treatments where there is a positive probability of an exogenously determined collusive price at 70 ECU first-order dominate the distribution observed in the baseline treatment, where such occurrences are precluded.[14] Note that the shift in distribution in not only driven by posted prices

[12] This data selection procedure is in accordance with our pre-study plan

[13] The distribution of purchase prices is almost identical, see appendix B.

[14] Kolmogorov-Smirnov tests for differences in distribution between $T0$ and the other treatments confirms that the price distributions in treatments $T10$, $T20$ and $T30$ are statistically different from

at 70 ECU. That is, if we condition the empirical price distributions on posted prices being less than 70 ECU, we still find that the price distributions from treatments $T10$, $T20$ and $T30$ first order dominates the price distribution from the treatment $T0$.[15]

We also find that, consistent with the theoretical predictions, the price distributions in treatments $T10$, $T20$ and $T30$ all have mass point at 70. The shares of posted price exactly at 70 ECU in $T10$, $T20$ and $T30$ respectively are 0.06, 0.1 and 0.39. The shares of posted price exactly at 70 ECU in $T0$, by comparison, is only 0.002. The differences in shares of posted prices at 70 ECU between $T0$ and the remaining treatments are all statistically significant.[16]

Figure 2: Price distribution by treatment

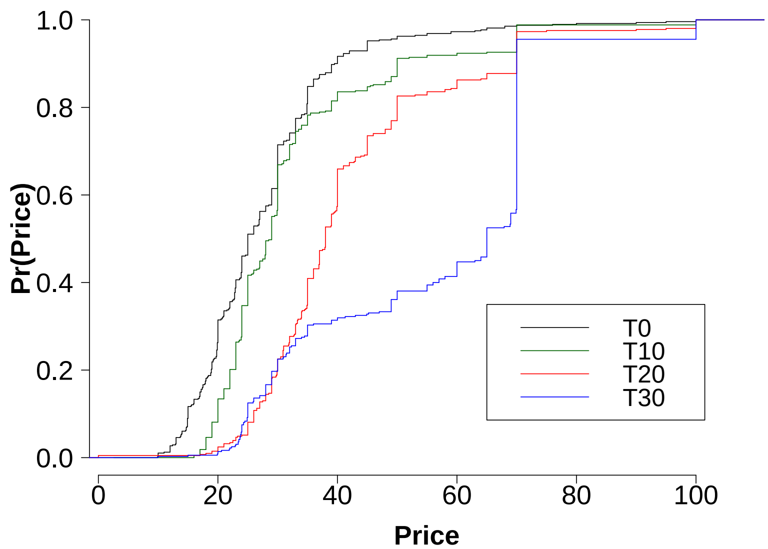

Note: First 20 games and games were the price was set exogenously are excluded.

Figure 3 shows the average posted prices by treatment. The average posted price

the price distribution in $T0$. All tests produce p-values below 0.001.

[15] Kolmogorov-Smirnov tests for differences in distribution between $T0$ and the other treatments, when these distribution condition on posted prices being less than 70 ECU, confirms that the price distributions in treatments $T10$, $T20$ and $T30$ are statistically different from the price distribution in $T0$. All tests produce p-values below 0.001.

[16] Wilcoxon rank-sum tests result in the following p-values: $T0$ vs. $T10$: $p = 0.003$; $T0$ vs. $T20$: $p = 0.004$; $T0$ vs. $T30$: $p = 0.004$

in treatment $T0$ is 27.6 ECU, which is close to the theoretical prediction of 24.4 ECU. The average prices in treatments $T10$, $T20$ and $T30$ are 32.3, 41.8 and 55.5 ECU respectively. The average prices in treatment $T20$ is close to the theoretical prediction of 42.6 ECU. Comparing posted prices across treatment we note that the differences between $T0$ and $T20$ and between $T0$ and $T30$ are statistically significant, while the difference between $T0$ and $T10$ is not.[17] We also note that average prices are significantly different across treatments $T10$, $T20$ and $T30$, with exeption of $T20$ vs. $T30$.[18]

We interpret these results as broadly supportive of our main theoretical prediction: sellers set higher prices when sellers are uncertain about whether collusion is going on compared to when collusion cannot happen. However, the extent of these higher prices seems to depend on the size of the probability of collusion. In particular, while prices in treatments $T0$ and $T20$ seem remarkably close to equilibrium, prices in $T10$ are lower than predicted and prices in $T30$ are higher than predicted.

[17] Wilcoxon rank-sum tests result in the following p-values: $T0$ vs. $T10$: $p = 0.394$; $T0$ vs. $T20$: $p = 0.009$; $T0$ vs. $T30$: $p = 0.004$

[18] See table 15 in appendix B for a table over all treatment by treatment comparisons tests

Figure 3: Average price by treatment

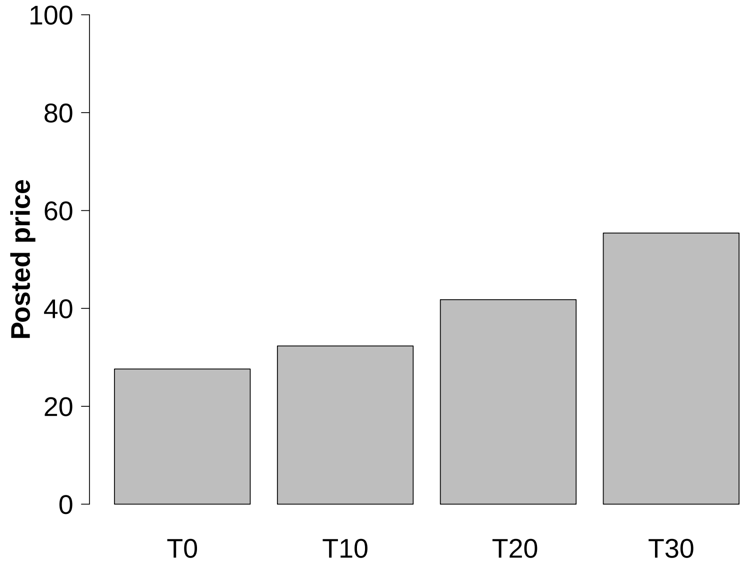

Note: First 20 games and games were the price was set exogenously are excluded. We exclude purchase price for informed buyers.

Focusing on equilibrium play, tables 3 and 4 present the share of price proposals within different ranges. Considering first treatment $T0$, we note that 86% of price proposals fall within the support of the equilibrium price distribution ($[14.8, 44.4]$). For the remaining treatments, the majority of price proposals also fall within the equilibrium range ($[26.5, 62.6] \cup 70$). The respective shares of price proposals within the equilibrium range in treatments $T10$, $T20$ and $T30$ are 55%, 85% and 70%.

That prices are set outside the equilibrium range, in particular in treatments $T10$ and $T30$, may of course be connected to out of equilibrium search. We will come back to this in section D, after we have analysed search behavior.

Table 3: Share of price proposals in different ranges in treatment $T0$

Treatment	$p \in [0, 14.8)$	$p \in [14.8, 44.4]$	$p \in (44.4, 100]$
$T0$	0.07	0.86	0.07

Note: First 20 games and games were the price was set exogenously are excluded.

Table 4: Share of price proposals in different ranges in treatments $T10$, $T20$ and $T30$

Treatment	$p \in [0, 26.5)$	$p \in [26.5, 62.6]$	$p \in (62.6, 70)$	$p = 70$	$p \in (70, 100]$
$T10$	0.43	0.49	0.00	0.06	0.01
$T20$	0.11	0.75	0.02	0.10	0.03
$T30$	0.14	0.31	0.12	0.39	0.04

Note: First 20 games and games were the price was set exogenously are excluded.

4.2 Search

Figure 4 presents average search rates by treatment. Search rates are low in all treatments, with rates of 0.06, 0.1, 0.05 and 0.12 in treatments $T0$, $T10$, $T20$ and $T30$, respectively. Although buyers search little on average, this represents a deviation from the theoretical no-search prediction. The differences in search behavior across treatments are not statistically significant.[19].

Figure 4: Share of buyers who search across treatments

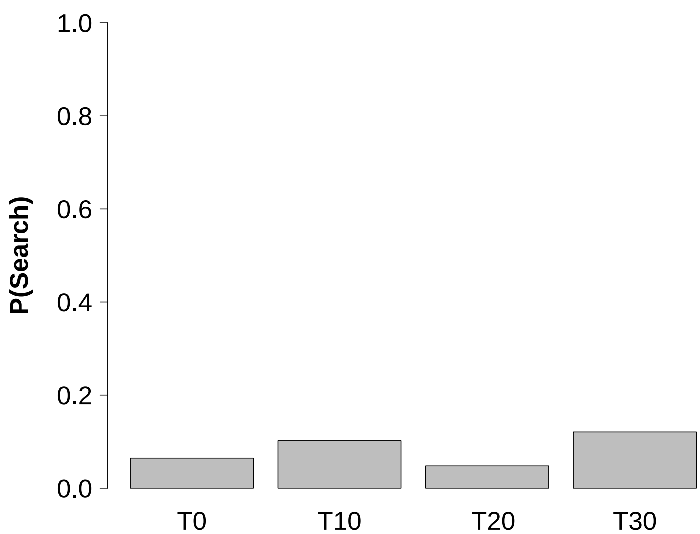

Note: First 20 games are excluded.

Tables 5 and 6 present the empirical search frequencies for different price ranges. We first note that search never takes place for price proposals below the support

[19]See table 17 in appendix B.

of the equilibrium price proposals. For the reaming price ranges, search decisions deviate from equilibrium predictions in some cases. Considering treatment $T0$, we note that buyers search to little for off path price proposals: Only a share of 0.65 decide to search when observing a price in excess of the upper bound of the equilibrium price distribution. In treatment $T10$ we may note that that a high share (0.28) decide to search following a price proposal of 70 ECU. Moreover, a higher share than in the other treatments search when observing a price in the equilibrium range $p \in [26.5, 62.6]$. This behavior is complemented by a large number of price proposals below the equilibrium range.

In treatment $T20$, search mainly takes place after having observed a price proposal of 70 ECU, though arguably with a low share. In treatment $T30$, however, this share is larger, and is, in combination with the large number of price proposals at 70 ECU, the main driver for the relatively large search rate observed in figure 4 for this treatment.

Table 5: Search by price range in treatment $T0$

Treatment	$p \in [0, 14.8)$	$p \in [14.8, 44.4]$	$p \in (44.4, 100]$
$T0$	0.0 (33)	0.02 (413)	0.65 (34)

Note: First 20 games are excluded. Number of price proposals in price range in parentheses.

Table 6: Search by price range in treatments $T10$, $T20$ and $T30$

Treatment	$p \in [0, 26.5)$	$p \in [26.5, 62.6]$	$p \in (62.6, 70)$	$p = 70$	$p \in (70, 100]$
$T10$	0 (185)	0.1 (214)	1 (1)	0.28 (27)	1 (5)
$T20$	0 (44)	0.01 (308)	0.33 (6)	0.08 (39)	1 (11)
$T30$	0 (49)	0.03 (112)	0.12 (43)	0.14 (140)	0.94 (16)

Note: First 20 games are excluded. Number of price proposals in price range in parentheses.

4.3 Implied reservation prices and optimal search behavior

Given the empirical price distribution we calculate the implied reservation price which we then use to determine the optimal search behavior in each treatment. We then compare the optimal search behavior to the actual search behavior to determine the degree to which buyers in our experiment best respond.

For treatment $T0$ the implied reservation price is simply the average posted price plus the search cost. For treatments $T10$, $T20$ and $T30$ we need to take into account that the observed price affects posterior beliefs about the price distribution. We do so by computing posterior beliefs that price has been set exogenously using the observed frequencies of price proposals at 70 ECU. The implied reservation prices are presented in table 7.

Table 7: Reservation prices implied by the empirical price distributions

Treatment	Implied reservation price
$T0$	47.6
$T10$	52.3/76.4
$T20$	61.8/79.8
$T30$	75.4/84.7

Note: First 20 games are excluded. For treatments $T10$, $T20$ and $T30$ the first number is the implied reservation price after observing a price different from 70 ECU and the second number is the implied reservation price when observing a price of 70 ECU.

If buyers in our experiment are best responding, given the empirical distribution of prices, they would always search whenever the price was in excess of the implied reservation prices in table 7, and never search when the price was below implied reservation prices in table 7. Comparing this optimal search behavior to the actual search behavior in table 8 we note that the overall pattern is remarkably consistent with buyers best responding. In particular, across all treatments more 90 percent of the decisions made by buyers are consistent with a best response. However, subjects search too much when observing prices below the implied reservation price and too little when the price is above the implied reservation price.

Table 8: Shares who search for prices at or below and above reservation price

Treatment	Observed price at or below implied reservation price	Observed price above implied reservation price	Overall share who best respond
$T0$	0.02	0.9	0.97
$T10$	0.08	1	0.92
$T20$	0.02	0.76	0.97
$T30$	0.09	0.94	0.91

Note: First 20 games are excluded.

4.4 Profits and optimal pricing strategy

The analysis in the previous section reveals that buyers to large degree best-respond given the empirical distribution of prices. In this section we consider whether sellers best respond given the search behavior of buyers, as well as the behavior of other sellers. To do so we estimate expected profits as function of posted price using the empirical price distribution and estimates of search propensities.[20] Estimated profits for all four treatments are presented in figure 5.

Consider first treatment $T0$. Note that profits are generally highest in equilibrium price range, $[14.8, 44.4]$. Comparing the profits for this price range with the frequency of posted prices within said ranges in table 3, we note that the large majority of posted prices (86 percent) are within the range that offers the highest profit. Thus, it appears that sellers in treatment $T0$ to large degree best respond in their pricing.

Consider next treatments $T10$, $T20$ and $T30$. First, note that profits for prices between 20 and 50 ECU and a price of 70 ECU in $T10$ are relatively high and almost identical. This price range accounts for 84 percent of posted prices in $T10$. In treatment $T20$, the price range between 25 and 70 ECU offers relatively high profits, and this range accounts for 89 percent of posted price. Finally, in treatment $T30$, the price range between 50 and 70 ECU offers relatively high profits. This range account for 58 percent of posted prices. Thus, in treatments $T10$, $T20$ and $T30$, seller behavior is also to a large degree consistent with with playing a best response.

[20] For treatment $T0$ search propensities for the entire range are estimated using a Gaussian kernel smoother with a bin with of 5. For treatments $T10$, $T20$ and $T30$, search propensities for prices in the range $[0, 70)$ are estimated using a Gaussian kernel smoother with a bin with of 5; for price proposals at 70 ECU, the search propensity is set to equal to the rates in table 6; for price proposals above 70 ECU, the search propensity is set to 1.

Figure 5: Price distribution by treatment

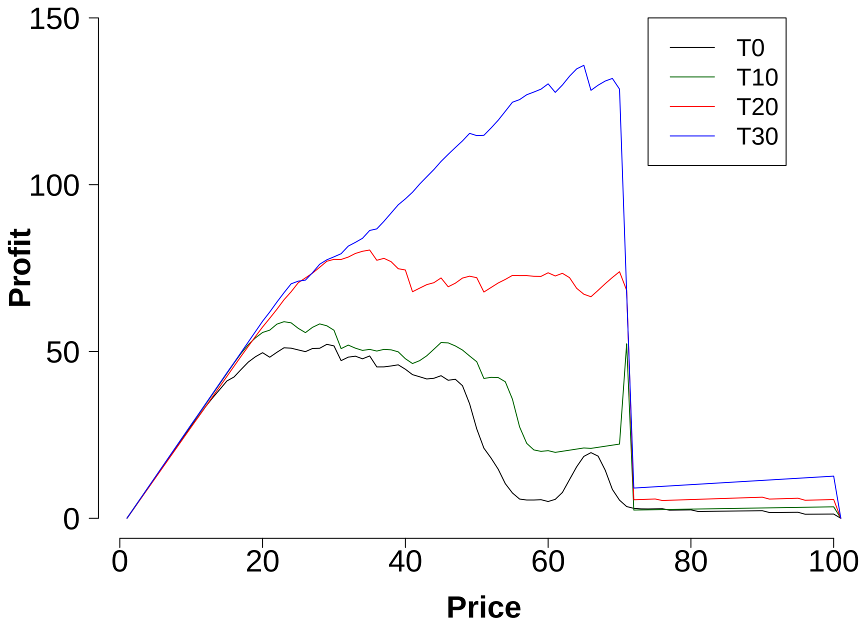

Note: Profit is calculated using the empirical price distribution from the last 20 rounds, excluding rounds were price is set exogenously, along with estimated search propensities.

5 A model of payoff-based learning

In payoff based learning models subjects are more prone to make choices associated with successful past outcomes than choices associated with less successful outcomes. We employ the much used reinforcement learning model developed by Roth and Erev (1995) and Erev and Roth (1998), which is much used to explain learning dynamics in experiments.[21] The model of reinforcement learning assumes that each action in a round of the game is associated with an *attraction*. The attraction of an action then determines the probability at which the action is played. The attraction of a given action depends on historical payoffs associated with that action along with an initial attraction. The initial attraction is estimated along with subjects sensitivity to

[21]See Masiliūnas (2023) for an recent overview of payoff learning models. Feltovich (2000) and Charness and Levin (2005) contrast bayesian updating with (reinforcement) learning. See also and Alós-Ferrer and Garagnani (2023) for an overview of this literature on Bayes vs learning.

attraction with respects to actions and tendency to forget. Estimates together with actual play can then be used to predict choices.[22]

The predicted shares of price proposals for the different ranges are presented in table 9. The predicted shares match some important features of the data, which are also consistent with the equilibrium predictions, specifically that a low shares of price proposals fall within the ranges $(62.6, 70)$ and $(70, 100]$. More important, the predicted shares match some important features of the data which the theory fails to explain. In particular, in our experiment a substantially higher share of price proposals fall within the range $[0, 26.5)$ in $T10$ than in $T20$ and $T30$. The predicted shares from our estimated learning model are quantitatively consistent with this pattern. Furthermore, in our experiment there seems to be a relationship between the probability at which the price is set exogenously and share of price proposals at 70 ECU. The learning model is also able to capture this feature of the data.

Table 9: Predicted share of price proposals in different ranges in treatments $T10$, $T20$ and $T30$

Treatment	$p \in [0, 26.5)$	$p \in [26.5, 62.6]$	$p \in (62.6, 70)$	$p = 70$	$p \in (70, 100]$
$T10$	0.17	0.54	0.05	0.19	0.05
$T20$	0.04	0.70	0.03	0.20	0.03
$T30$	0.04	0.28	0.07	0.58	0.03

Note: First 20 games are excluded. Estimated on pooled data.

The predicted search propensities for price proposals within the different ranges are presented in table 10. The predicted search behaviour matches some important features of the data, which are also qualitatively consistent with the equilibrium predictions. In particular, the learning model predicts no search for prices in the lowest range and a high search propensity for prices in the highest range. Furthermore, the learning model predicts low search propensities for price proposals in the range $[26.5, 62.6)$ and higher search propensities for prices in the rage $(62.6, 70)$ than for prices at 70 ECU. Note finally that the search behaviour matches some important

[22] In appendix C we describe the details of the the model of reinforcement learning as well as the estimation procedure.

features of the data which the theory fails to explain. Specifically, the learning model predicts a higher search propensity for price proposals at 70 ECU in treatment $T10$ than in $T20$ and $T30$.

Table 10: Predicted search propensity for price proposals in different ranges in treatments $T10$, $T20$ and $T30$

Treatment	$p \in [0, 26.5)$	$p \in [26.5, 62.6]$	$p \in (62.6, 70)$	$p = 70$	$p \in (70, 100]$
$T10$	0.00	0.02	0.34	0.27	0.82
$T20$	0.00	0.01	0.33	0.19	0.82
$T30$	0.00	0.06	0.30	0.13	0.82

Note: First 20 games are excluded. Estimated on pooled data.

6 Conclusion

In this paper, we explore the impact of potential collusion on equilibrium prices, both theoretically and experimentally. The theoretical model suggests that fear of collusion affects consumer search behavior, reducing the incentive to search further after encountering a high price. This behavior allows non-colluding sellers to set higher prices. In the laboratory experiment, we vary the probability of collusion within treatments. Qualitatively, our results align reasonably well with the theoretical predictions, showing an increase in average prices with potential collusion. Despite some deviations from theoretical expectations, both buyers and sellers demonstrate strong best-response behavior. The discrepancies are to some extent attributed to subjects' sensitivity to experience, as explored through a reinforcement learning model.

References

Abrams, E., M. Sefton, and A. Yavas (2000): "An experimental comparison of two search models," *Economic Theory*, 16, 735–749.

Alós-Ferrer, C. and M. Garagnani (2023): "Part-time bayesians: Incentives and behavioral heterogeneity in belief updating," *Management Science*.

Atayev, A. (2022): "Uncertain product availability in search markets," *Journal of Economic Theory*, 204, 105524.

Ball, S. B. and P.-A. Cech (1996): "Subject pool choice and treatment effects in economic laboratory research," *Research in experimental economics*, 6, 239–292.

Bellemare, C., L. Bissonnette, and S. Kröger (2016): "Simulating power of economic experiments: the powerBBK package," *Journal of the Economic Science Association*, 2, 157–168.

Benabou, R. and R. Gertner (1993): "Search with learning from prices: does increased inflationary uncertainty lead to higher markups?" *The Review of Economic Studies*, 60, 69–93.

Burdett, K. and K. L. Judd (1983): "Equilibrium Price Dispersion," *Econometrica*, 51, 955–969.

Campbell, C., G. Ray, and W. A. Muhanna (2005): "Search and collusion in electronic markets," *Management Science*, 51, 497–507.

Cason, T. N. and S. Datta (2006): "An experimental study of price dispersion in an optimal search model with advertising," *International Journal of Industrial Organization*, 24, 639–665.

Cason, T. N., D. Friedman, and E. Hopkins (2021): "An experimental investigation of price dispersion and cycles," *Journal of Political Economy*, 129, 789–841.

Cason, T. N., D. Friedman, et al. (2003): "Buyer search and price dispersion: a laboratory study," *Journal of Economic Theory*, 112, 232–260.

Cason, T. N. and S. D. Mago (2010): "Costly buyer search in laboratory markets with seller advertising," *The Journal of Industrial Economics*, 58, 424–449.

Charness, G. and D. Levin (2005): "When optimal choices feel wrong: A laboratory study of Bayesian updating, complexity, and affect," *American Economic Review*, 95, 1300–1309.

Dana Jr, J. D. (1994): "Learning in an equilibrium search model," *International Economic Review*, 745–771.

Davis, D. D. and C. A. Holt (1996): "Consumer search costs and market performance," *Economic Inquiry*, 34, 133–151.

Diamond, P. A. (1971): "A model of price adjustment," *Journal of economic theory*, 3, 156–168.

Erev, I. and A. E. Roth (1998): "Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria," *American economic review*, 848–881.

Feltovich, N. (2000): "Reinforcement-based vs. belief-based learning models in experimental asymmetric-information games," *Econometrica*, 68, 605–641.

Fischbacher, U. (2007): "z-Tree: Zurich toolbox for ready-made economic experiments," *Experimental economics*, 10, 171–178.

Fréchette, G. R. (2016): "Experimental economics across subject populations," *The handbook of experimental economics*, 2, 435–480.

Greiner, B. (2015): "Subject pool recruitment procedures: organizing experiments with ORSEE," *Journal of the Economic Science Association*, 1, 114–125.

Grether, D. M., A. Schwartz, and L. L. Wilde (1988): "Uncertainty and shopping behaviour: an experimental analysis," *The Review of Economic Studies*, 55, 323–342.

Harrington Jr, J. E., R. H. Gonzalez, and P. Kujal (2016): "The relative efficacy of price announcements and express communication for collusion: Experimental findings," *Journal of Economic Behavior & Organization*, 128, 251–264.

Heggedal, T.-R., L. Helland, and E. R. Moen (2023): "Sequential Price Setting: Theory and Evidence from a Lab Experiment," *International Economic Review*, x, x.

Helland, L., E. R. Moen, and E. Preugschat (2017): "Information and coordination frictions in experimental posted offer markets," *Journal of Economic Theory*, 167, 53–74.

Hong, H. and M. Shum (2006): "Using price distributions to estimate search costs," *The RAND Journal of Economics*, 37, 257–275.

Janssen, M., P. Pichler, and S. Weidenholzer (2011): "Oligopolistic markets with sequential search and production cost uncertainty," *The RAND Journal of Economics*, 42, 444–470.

Janssen, M. and S. Shelegia (2015): "Consumer search and double marginalization," *American Economic Review*, 105, 1683–1710.

——— (2020): "Beliefs and consumer search in a vertical industry," *Journal of the European Economic Association*, 18, 2359–2393.

Janssen, M. C. (2020): "Vertical contracts in search markets," *International Journal of Industrial Organization*, 70, 102527.

Janssen, M. C., A. Parakhonyak, and A. Parakhonyak (2017): "Non-reservation price equilibria and consumer search," *Journal of Economic Theory*, 172, 120–162.

Kahneman, D. and A. Tversky (1972): "Subjective probability: A judgment of representativeness," *Cognitive psychology*, 3, 430–454.

Lauermann, S., W. Merzyn, and G. Virág (2012): "Learning and Price Discovery in a Search Model," Tech. rep., Working Paper.

Lester, B. (2011): "Information and prices with capacity constraints," *American Economic Review*, 101, 1591–1600.

Lubensky, D. (2017): "A model of recommended retail prices," *The RAND Journal of Economics*, 48, 358–386.

Masiliūnas, A. (2023): "Learning in rent-seeking contests with payoff risk and foregone payoff information," *Games and Economic Behavior*, 140, 50–72.

Mauring, E. (2017): "Learning from trades," *The Economic Journal*, 127, 827–872.

——— (2020): "Informational cycles in search markets," *American Economic Journal: Microeconomics*, 12, 170–192.

Moellers, C., T. Stühmeier, and T. Wenzel (2016): "Search costs in concentrated markets: An experimental analysis," Tech. rep., DICE Discussion Paper.

Montag, F. and C. Winter (2020): "Price transparency against market power," *Available at SSRN 3256476*.

Moraga-González, J. L., Z. Sándor, and M. R. Wildenbeest (2023): "Consumer search and prices in the automobile market," *The Review of Economic Studies*, 90, 1394–1440.

Morgan, J., H. Orzen, and M. Sefton (2006): "An experimental study of price dispersion," *Games and Economic Behavior*, 54, 134–158.

Nilsson, A. et al. (1999): "Transparency and competition," Tech. rep., mimeo, Stockholm School of Economics.

Nishida, M. and M. Remer (2018): "Lowering consumer search costs can lead to higher prices," *Economics Letters*, 162, 1–4.

Obradovits, M. and P. Plaickner (2020): "Price-Directed Search and Collusion," Tech. rep., Working Papers in Economics and Statistics.

Orzen, H. (2008): "Counterintuitive number effects in experimental oligopolies," *Experimental Economics*, 11, 390–401.

Ouwersloot, H., P. Nijkamp, and P. Rietveld (1998): "Errors in probability updating behaviour: Measurement and impact analysis," *Journal of Economic Psychology*, 19, 535–563.

Petrikaitė, V. (2016): "Collusion with costly consumer search," *International Journal of Industrial Organization*, 44, 1–10.

Potters, J. and S. Suetens (2013): "Oligopoly experiments in the current millennium," *Journal of Economic Surveys*, 27, 439–460.

Robert, J. and D. O. Stahl (1993): "Informative price advertising in a sequential search model," *Econometrica: Journal of the Econometric Society*, 657–686.

Roth, A. E. and I. Erev (1995): "Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term," *Games and economic behavior*, 8, 164–212.

Schultz, C. (2005): "Transparency on the consumer side and tacit collusion," *European Economic Review*, 49, 279–297.

——— (2017): "Collusion in markets with imperfect price information on both sides," *Review of industrial organization*, 50, 287–301.

Shadarevian, V. (2022): "Collusive pricing with consumer search," *Available at SSRN 4255948.*

Snowberg, E. and L. Yariv (2021): "Testing the waters: Behavior across participant pools," *American Economic Review*, 111, 687–719.

Stahl, D. (1989): "Oligopolistic Pricing with Sequential Consumer Search," *American Economic Review*, 79, 700–712.

Tappata, M. (2009): "Rockets and feathers: Understanding asymmetric pricing," *The RAND Journal of Economics*, 40, 673–687.

Tversky, A. and D. Kahneman (1971): "Belief in the law of small numbers." *Psychological bulletin*, 76, 105.

Varian, H. R. (1980): "A Model of Sales," *The American Economic Review*, 70, 651–659.

Wolinsky, A. (1986): "True monopolistic competition as a result of imperfect information," *The Quarterly Journal of Economics*, 101, 493–511.

Yang, H. and L. Ye (2008): "Search with learning: understanding asymmetric price adjustments," *The RAND Journal of Economics*, 39, 547–564.

A Model extension

A.1 Collusion price less than the Stahl price

Suppose first that the collusion price is less than the Stahl price p^s.Then the following holds

Result 4. *Suppose $p^M \leq p^s$. Then the Stahl equilibrium, with no mass point at p^M, and no consumer search at the support $[p_0, p_1]$, $p_1 = p^s$ is an equilibrium.*

Suppose the other seller (and the buyers) play according to Stahl. Then the seller in question is indifferent between setting any price in $[p_0, p_1]$, including setting p^M. Hence playing Stahl is a best response. Consider then the consumer. Since all prices are at or below p^s, the consumer has no incentives to search if when observing a price at or below p_1. In particular, the consumer has no incentives to search when $p = p^m$.

Note also that (in the absence of collusion) there cannot be an equilibrium in which there is a mass point at p^M, as the usual undercutting argument applies: If there is a mass point at p^M, the best response of a seller is to undercut and set the price just below p^M. Hence the strategy is not a best response and we cannot be in equilibrium.

A.2 Support of p^M above p^s

We retain the assumption that firms collude with probability x. However, we assume that the colluding price is stochastic with distribution $G(p)$. For the moment, we assume that the distribution is continuous with support $(p^{\min}, p^{\max})$, $p^{\min} \geq 0$, $p^{\max} \leq 1$.

We first consider an equilibrium candidate in which 1) sellers randomise on two intervals $I^N = [p_0, p_1]$ and $I^M = [p^{\min}, p^{\max}]$. We first assume that $p^{\min}$ is sufficiently high so that the two distributions do not overlap (we will come back to this later). We denote the distribution given that the seller is setting a price in I^N by F^N, and the distribution given that it sets a price in I^M by F^M. Let ρ denote the probability that a seller sets a price in I^M. For all $p \in I^M$, let $q(p) \in [0,1]$ denote the probability that the buyer searches, and let $Q(p) = \int_p^{p^{\max}} q(\tau) dF^M(\tau)$. Let $\bar{Q} = Q(p^{\min})$. The profit of a firm setting $p \in F^N$ is given by

$$\begin{aligned} \pi(p) &= p\left(u + \rho\bar{Q} + \rho I + (1-\rho)(1 - F^N(p))\right) \\ &= p(\tilde{u} + \tilde{I}) \end{aligned} \quad (11)$$

where $\tilde{u} = (1 + \rho Q)u + \rho I$ and $\tilde{I} = (1-\rho)I$

Equilibrium is a value $\rho \in [0,1]$, two distribution functions $F^N(p)$ and $F^M(p)$ with support I^N and I^M, respectively, and a function $q(p)$ on I^M satisfying the following conditions:

1. Equal profit when setting p^M and p_1:

2. Equal profits for all $p \in [p_0, p_1]$ implying that $F(p)$ given by

$$1 - F(p) = \frac{\tilde{u}(p_1 - p)}{p\tilde{I}} \quad (12)$$

   where p_0 is defined by $F(p_0) = 0$.

3. All prices in the support of F^M give the same profits to seller, and this profit is equal to the profit if setting a price in I^N.

4. Uninformed consumers are indifferent between searching or not searching at p_1:

$$(1-\rho)\int_{p_0}^{p_1} F(p)dp = c \quad (13)$$

5. When observing a price $p \in I^M$, uninformed consumers are 1) either indifferent between searching and not searching (if $q(p) < 1$) or prefer not to search (if $q = 1$).

Let us first make some observation. The first thing to note is that F^M has no mass points. Suppose it had a mass point at p'. Then since G is without mass points, the posterior probability that there is collusion is 0, and the consumers will search with probability 1. But this implies that the profit if setting p' is very low, a contradiction (to be sharpen).

the consumers know with probathere exists no equilibrium in which the consumers do not randomise for any price in I^M. Suppose it did, i.e., suppose that the consumers randomise for all prices in I^M, and particularly for $p^{\min}$. But then sellers would be strictly better off setting $p^{\min}$ than setting p^1, a contradiction.

B Results

B.1 Inference: Last 20 games

Table 11: Wilcoxon rank-sum test (p-values): Posted price

	$T0$	$T10$	$T20$	$T30$
$T0$				
$T10$	0.394			
$T20$	0.009	0.026		
$T30$	0.004	0.009	0.132	

Note: First 20 games and games were the price was set exogenously are excluded.

Table 12: Wilcoxon rank-sum test (p-values): Share of posted prices at 70 ECU

	$T0$	$T10$	$T20$	$T30$
$T0$				
$T10$	0.003			
$T20$	0.004	0.090		
$T30$	0.004	0.012	0.009	

Note: First 20 games and games were the price was set exogenously are excluded.

Table 13: Wilcoxon rank-sum test (p-values): Share who search.

	$T0$	$T10$	$T20$	$T30$
$T0$				
$T10$	0.406			
$T20$	0.256	0.091		
$T30$	0.101	0.809	0.075	

Note: First 20 games are excluded.

B.2 Results and inference: All games

Table 14: Treatment measures by treatment

	Average price	Share of posted prices at 70 ECU	Share who search
$T0$	32.6	0.01	0.091
$T10$	37.3	0.8	0.135
$T20$	45.7	0.10	0.097
$T30$	53.7	0.31	0.135

Note: Games were the price was set exogenously are excluded from average price and share of posted prices at 70 ECU.

Table 15: Wilcoxon rank-sum test (p-values): Posted price

	$T0$	$T10$	$T20$	$T30$
$T0$				
$T10$	0.310			
$T20$	0.004	0.064		
$T30$	0.002	0.041	0.485	

Note: Games were the price was set exogenously are excluded.

Table 16: Wilcoxon rank-sum test (p-values): Share of posted prices at 70 ECU

	$T0$	$T10$	$T20$	$T30$
$T0$				
$T10$	0.012			
$T20$	0.005	0.172		
$T30$	0.005	0.013	0.012	

Note: First 20 games and games were the price was set exogenously are excluded.

Table 17: Wilcoxon rank-sum test (p-values): Share who search.

	$T0$	$T10$	$T20$	$T30$
$T0$				
$T10$	0.065			
$T20$	0.520	0.228		
$T30$	0.077	0.699	0.149	

Note: First 20 games are excluded.

B.3 Results by round

Figure 6: Average price by game and treatment

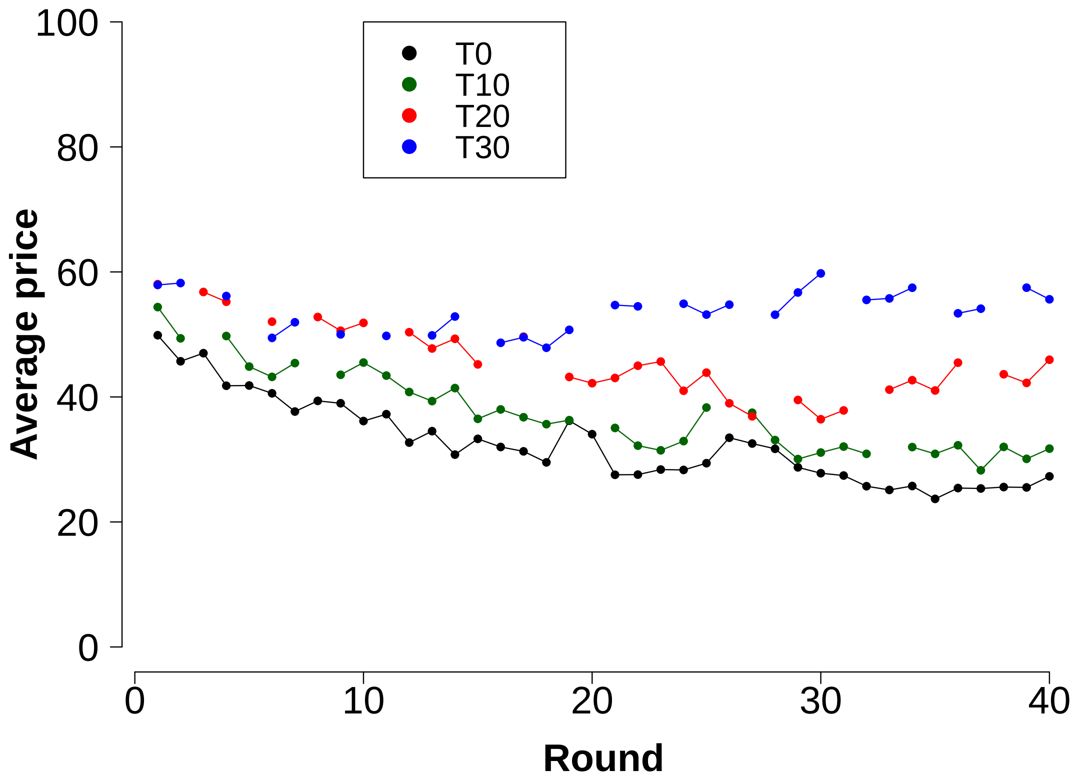

Note: Games were the price was set exogenously are excluded.

Figure 7: Search behavior by game and treatment

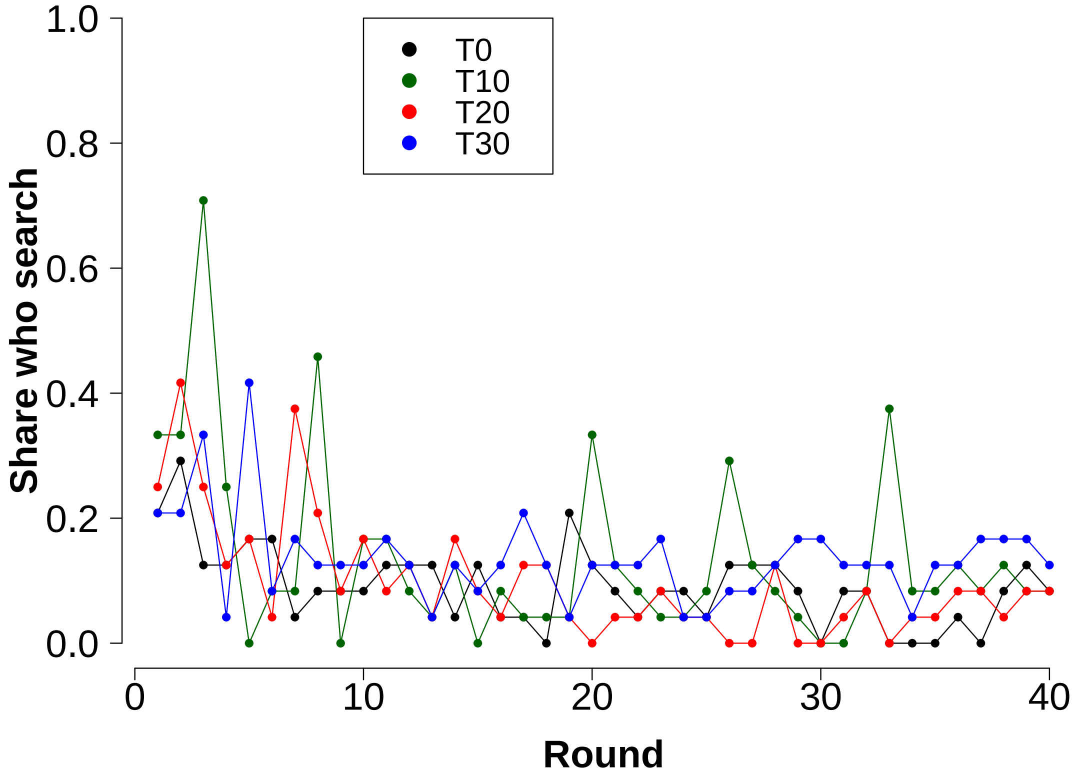

B.4 Other results

C Learning model

The model of reinforcement learning assumes that each action in a round of the game is associated with an *attraction*. The attraction of an action then determines the probability at which the action is played. For sellers, we let $A_i^{Seller}(t)$ denote the attraction of posting a price in the range P_i in round t. The attraction for sellers associated with a price range P_i is updated according to the following rule

$$A_i^{Seller}(t) = \begin{cases} \phi_{Seller} A_i^{Seller}(t-1) + \pi(t) & \text{if } p(t) \in P_i \\ \phi_{Seller} A_i^{Seller}(t-1) & \text{otherwise} \end{cases} \tag{14}$$

where $\pi(t)$ is the realised payoff in round t and ϕ_{Seller} is the recency parameter. The probability of posting a price in range P_j is then

$$Pr_i(t) = \frac{\exp(\lambda_{Seller} A_j^{Seller}(t-1))}{\sum \exp(\lambda_{Seller} A_i^{Seller}(t-1))} \tag{15}$$

where λ_{Seller} represents sensitivity to attraction.

For buyers, let $s_i(t) \in \{Buy, Search\}$ denote the action of a buyer when observing a price in range P_i in round t. We let $A_{i,s}^{Buyer}(t)$ denote the attraction of s when observing a price in the range P_i in round t. The attraction for buyers is then updating as follows

$$A_{i,s}^{Buyer}(t) = \begin{cases} \phi_{Buyer} A_{i,s}^{Buyer}(t-1) + \pi(t) & \text{if } p(t) \in P_i \text{ and } s_i(t) = s \\ \phi_{Buyer} A_{i,s}^{Buyer}(t-1) & \text{otherwise} \end{cases} \tag{16}$$

The probability of s when observing a price in range P_i is then

$$Pr_{i,s}(t) = \frac{\exp(\lambda_{Buyer} A_{i,s}^{Buyer}(t-1))}{\exp(\lambda_{Buyer} A_{i,Search}^{Buyer}(t-1)) + \exp(\lambda_{Buyer} A_{i,Buy}^{Buyer}(t-1))} \tag{17}$$

When fitting the model we consider five different price ranges: $[0, 26.5)$, $[26.5, 62.6]$, $(62.6, 70)$, 70 and $(70, 100]$.[23] We estimate initial attraction, $A_i^{Seller}(0)$ and $A_{i,s}^{Buyer}(0)$, ϕ_{Seller}, ϕ_{Buyer}, λ_{Seller} and λ_{Buyer}, but $A_i^{Seller}(0)$ is normalised to zero for the lowest price range ($[0, 26.5)$), and $A_{i,Buy}^{Buyer}(0)$ is normalised to zero for all price ranges.

The model estimates are presented in tables 18a and 18b. The estimation pools

[23] The model becomes difficult to estimate with more more granular ranges.

data from all treatments.[24] The recency parameter is close to 1 for both buyers and sellers which indicates that learning is persistent and that past and recent payoffs carry an almost equal weight in current decisions.

Table 18: Model estimates

Parameter	Estimate (Std. err.)
ϕ_{Seller}	0.943 (0.005)
λ_{Seller}	0.003 (<0.001)
$A_1^{Seller}(0)$	0 (normalised)
$A_2^{Seller}(0)$	842.2 (103.3)
$A_3^{Seller}(0)$	141.4 (109.3)
$A_4^{Seller}(0)$	-180.0 (105.8)
$A_5^{Seller}(0)$	-88.3 (127.7)

(a) Seller

Parameter	Estimate (Std. err.)
ϕ_{Buyer}	0.985 (0.006)
λ_{Buyer}	0.005 (0.001)
$A_{1,Search}^{Buyer}(0)$	-42066.2 (2.08e+08)
$A_{2,Search}^{Buyer}(0)$	-371.9 (76.8)
$A_{3,Search}^{Buyer}(0)$	-226.2 (67.2)
$A_{4,Search}^{Buyer}(0)$	-220.2 (46.0)
$A_{5,Search}^{Buyer}(0)$	478.2 (122.5)

(b) Buyer

D Optimal pricing strategy

In this section we consider whether sellers best respond given the search behavior of buyers, as well as the behavior of other sellers. To do so we estimate expected profits for posted prices in different ranges for all treatments and then compare to the frequency of posted prices within these ranges.

Consider first treatment $T0$. In figure 5 we present expected profits for three different price ranges. Note that the equilibrium price range stands out compared to the two others by offering a significantly higher expected profit: The average profit for posted prices in the range $[14.8, 44.4]$ is 52.5 ECU while it is 33 ECU for prices posted in the lower range ($[0, 14.8)$) and 22.6 ECU for prices posted in the higher range

[24]Given that play in initial rounds is very similar across treatments this seems well justified, see figures 6 and 7 in appendix B.

((44.4, 100]). Comparing the profits for the the three price ranges with the frequency of posted prices within said ranges in table 3, we note that the large majority of posted prices (86 percent) are within the the range that offers the highest profit. Thus, it appears that sellers in treatment $T0$ to large degree best respond in their pricing.

Figure 8: Expected profit by price range in treatment $T0$

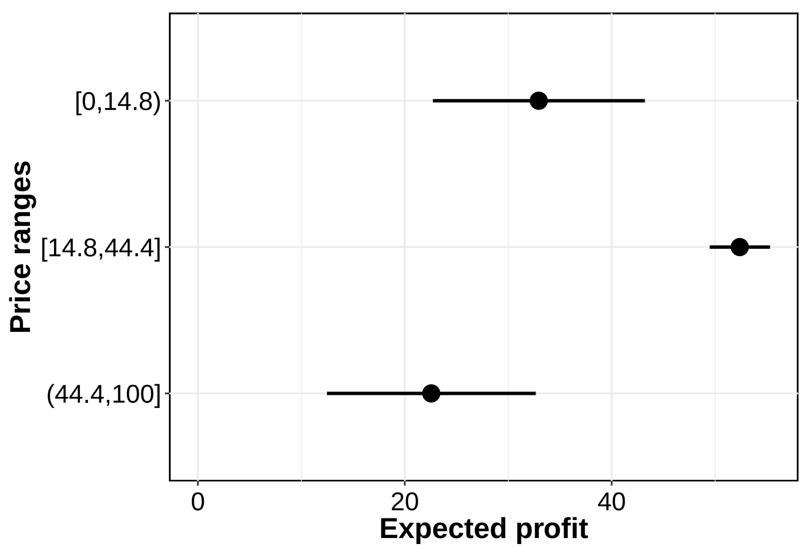

Note: First 20 games and games were the price was set exogenously are excluded. Expected profit (points) is inferred directly from å dummy regression. Lines indicate 95% confidence intervals.

Consider next treatments $T10$, $T20$ and $T30$. Note first that profits are generally higher in treatments $T10$, $T20$ and $T30$ than in treatment $T0$, and that among the treatments with potential collusion, the profits vary positively with the exogenous probability of collusion. Second, note that profits for the two lowest price ranges in $T10$ are almost identical. This may account for the large share of price proposals for the lowest range ([0, 26.5)) in $T10$.

Figure 9: Expected profit by price range in treatments $T10$, $T20$ and $T30$

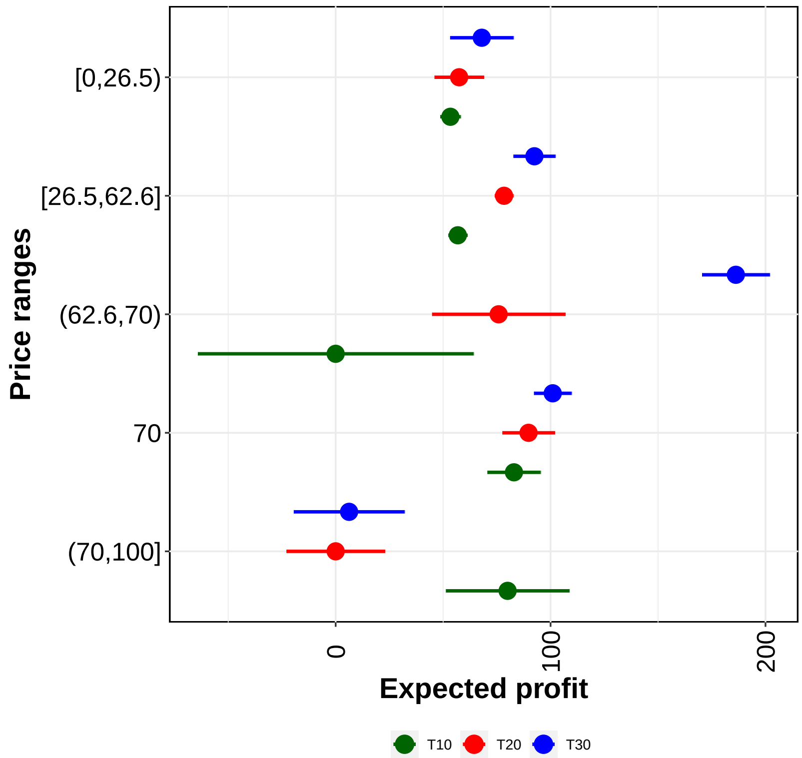

Note: First 20 games and games were the price was set exogenously are excluded.