



**Strategic Inattention
in Product Search**

**Adrian Hillenbrand
Svenja Hippel**





Strategic Inattention in Product Search

Adrian Hillenbrand / Svenja Hippel

November 2017,
revised September 2019

Strategic Inattention in Product Search*

Adrian Hillenbrand[†] and Svenja Hippel[‡]

September 12, 2019

Abstract

Rapid technological developments in online markets fundamentally change the relationship between consumers and sellers. Online platforms can now easily gather data about the consumer and his search behavior, that allow for price discrimination. Therefore the consumers' product search becomes a strategic choice. Consumers face a trade-off: Search intensely and receive a better fit at a potentially higher price or restrict search behavior – be strategically inattentive – and receive a worse fit, but maybe a better deal. We study the resulting strategic buyer-seller interaction theoretically as well as experimentally. Our experimental results shed a critical light on the added value for consumers through the rise of online platforms.

Keywords: strategic inattention, price discrimination, information transmission, consumer choice, experiment

JEL Classification: D11, D42, D82, D83, L11

*We would like to thank Phil Brookins, Christoph Engel, Lars Freund, Dominik Grafenhofer, Dorothea Kübler, Lisa Lenz, Daniel Martin, Anne-Katrin Roesler, Sebastian Schaube, Patrick Schmitz, Gyula Seres and Fabian Winter, as well as audiences at the Max Planck Institute for Research on Collective Goods, the Summer Meeting of Young Philosophers and Economists 2016, the 33rd annual NABE Economic Policy Conference, ESA Meetings in San Diego and Vienna, the 12th Nordic Conference in Gothenburg, the 11th M-BEES, EARIE 2018, and the annual meeting of the Verein for Socialpolitik for helpful comments and discussion. Financial support from the Max Planck Society is gratefully acknowledged. IRB approval due to a block approval by the University of Bonn.

[†]Max Planck Institute for Research on Collective Goods, Kurt-Schumacher-Str. 10, 53113 Bonn, Germany. E-mail: hillenbrand@coll.mpg.de.

[‡]University of Würzburg, Sanderring 2, 97070 Würzburg, Germany. E-mail: svenja.hippel@uni-wuerzburg.de.

1 Introduction

Rapid technological developments in online markets fundamentally change the relationship between consumers and sellers. The rise of online platforms increases the transparency for consumers in many markets because a multitude of products can now easily be accessed and browsed through on a single web page. At first sight, this is beneficial for consumers since they can find more relevant and better-fitting product offers. But online platforms can also more easily gather data about consumers, in particular about those with a more intense search behavior on the particular site. In the extreme, personalized search results based on real-time data then allow for perfect price discrimination. Increasingly, the consultancy industry champions this type of ‘behavioral pricing’ and technological advancements will further simplify information extraction.

Therefore, the extent of consumers’ product search becomes a strategic choice and consumers are potentially better off by restricting their search behavior. Throughout the paper, we will call this active decision to restrict product search *strategic inattention*.¹ Whether consumers appreciate the strategic situation and react suitably by being strategically inattentive and whether sellers indeed exploit the collected information might have a severe impact on consumer welfare and on the need for regulation in online markets. Given the rising importance of online platforms, it is important to understand how consumers adapt to these new markets.

Recent models in information economics provide theoretical insights into such situations of endogenous information transmission between a buyer and a seller. [Bergemann et al. \(2015\)](#) describe the set of possible market outcomes and thus the boundaries of price discrimination when the seller can segment the market based on additional information about the buyer’s taste. Given the different extents of such buyer-related information, [Carrasco et al. \(2018\)](#) derive the selling mechanism that maximizes revenue for the seller. [Roesler and Szentes \(2017\)](#) and [Condorelli and Szentes \(2018\)](#), on the other hand, analyze the situation from the buyer’s perspective when the seller optimally responds to the chosen information design. In the two papers, the seller learns the choice of the buyer before setting a price and the buyer-optimal decision is not to be informed perfectly. In [Roesler and Szentes \(2017\)](#), the buyer’s valuation

¹The term ‘strategic inattention’ is different from ‘rational inattention’ which is discussed in the literature, started by the seminal papers of [Sims \(2003, 2006\)](#). These papers deal with the unobservable psychological costs of processing information based on a Shannon entropy (see [Caplin et al. \(2017\)](#) for a behavioral characterization). Subjects are already shown by [Martin \(2016, 2017\)](#), theoretically as well as empirically, to behave as if they are rationally inattentive. In our setup, we exclude psychological as well as all exogenous search costs. All search costs are of purely strategic nature and are caused through the seller’s ability to extract information of the buyer’s valuation by observing the search behavior.

distribution is given and initially unknown to both sides. The buyer chooses the optimal precision of a costless signal about his valuation that enables her to improve his buying decision. In [Condorelli and Szentes \(2018\)](#), the buyer directly chooses his valuation distribution. In both cases, the information transmission generates a trade-off between a higher valuation and a higher price (see also [Ichihashi, 2019](#)).

In this paper, to the best of our knowledge, we provide a first experimental analysis of a model of buyer-optimal information design. The experimental setup is a stylized market that features the key component of information transmission: sellers can observe the search behavior of buyers before setting a price. We first analyze the situation theoretically. Our setup is thereby close to [Condorelli and Szentes \(2018\)](#) in that the buyer's choice directly determines his value distribution. However, in order to receive testable hypotheses for the experimental study, we impose a specific structure on the buyer's decision space. For the experiment we create a decision environment inspired by a familiar search situation on an online platform.

In principle, sellers have many options when and where to collect consumer search data. One example where information transmission would be possible, is filters on booking platforms. These filters can be used to restrict the offers displayed, e.g., to hotels that are close to the beach and come with a gym. The use of such provided search tools potentially transmits information about the consumer's valuation for the filtered products. Crucially, this data is available to the seller before the list of search results and the product prices are displayed to the consumer. Consequently, this information can be used to price-discriminate and to steer consumers towards buying more expensive products. Consumers then face a trade-off between product fit and price: They can search intensely for a well-fitting product but, due to the information transmission to the seller, might face a higher price. Thus, they are potentially better off being strategically inattentive to the available product information. While they might then receive a worse fit, they might end up with a lower price and a better deal.

First empirical market studies (see, e.g., [Mikians et al., 2012](#); [Hannak et al., 2014](#)) provide hints on the extent of the use of consumer data (see [Acquisti et al., 2016](#), for a more general overview over the literature). A recent study on online travel agents by [Hunold et al. \(2018\)](#) also showed that the ranking of search results on booking platforms is indeed not only oriented on displaying the best-fitting product for the consumer. These findings suggest the existence of price discrimination practices on online platforms, but they cannot evaluate how consumers react to them. With the volatility of prices in these markets, price discrimination and other

causes for price changes are hard to distinguish (Visser et al., 2014). Therefore, it is hard to tell which search strategy could be optimal in a given situation. Consequently, using field data alone is not sufficient to analyze causes for specific buyer search behavior and to understand whether they react suitably.

The controlled environment of the laboratory thus allows us to analyze the buyer, but also the seller behavior and the development of the interaction over time. In particular, we can fix which information sellers receive about buyers' search behavior. We can ensure that buyers are aware of this, such that the information transmission is unambiguously taking place and is common knowledge. In the experiment, sellers offer one product out of a given set of products with different characteristics. Buyers can make use of a search technology in the form of filters to restrict this product set beforehand by removing bad options.² Sellers observe the number of filters chosen and set a fixed price for all the remaining products. This allows them to condition their pricing behavior on the buyer's filter choice. Finally, buyers receive a random product from the remaining set and decide whether to buy it at the price set by the seller. While clearly a simplification of the decision making environment in the field, this setup captures the key features of information transmission in the markets we are interested in.

In the Nash equilibrium, product prices increase weakly in the buyer's filter choice and the buyer consequently restricts the number of filters used. The results of our experimental study are generally in line with this theoretical prediction. Sellers set higher prices for higher filter choices and the buyers are strategically inattentive, although not to the extent predicted by theory. Interestingly, it is only the sellers who profit from the buyer's use of the search technology: Their surplus increases with higher filter choices of the buyers. Buyers themselves though do not gain from a more extensive use of the search technology. At an individual level, we observe quite distinct, but recurring patterns of filter choice over time.

Our paper allows first behavioral insights into the growing market of online platforms and new challenges that might arise with further technological advancement. Our results shed a critical light on the added value for consumers through the rise of online platforms. They call the benefits of progressing personalization of online search results into question. In particular, our results suggest that it is the sellers and not the buyers who profit from the rapid technological developments in online markets.

²Bad options are purely based on (induced) taste over product characteristics. There are no quality differences between products in the experiment. Note also that the experimental setup does not allow for sequential search of any type as the buyer's decision is to choose a number of filters.

Related literature. The theoretical literature on behavior-based price discrimination discusses multi-period models that capture situations where a recurring customer can be identified, for example, through the use of browser cookies. [Fudenberg and Villas-Boas \(2006\)](#) provide an overview over these models, in which the seller can charge different prices based on past buying decisions of consumers (see also [Hart and Tirole, 1988](#); [Villas-Boas, 2004](#)). Our model goes beyond this in capturing the use of search data that reveal information about the consumer’s valuation in real time. [Dengler and Prüfer \(2018\)](#) and [Montes et al. \(2018\)](#) also have a theoretical setup where the consumer can decide between an anonymous channel and one where all personal data are revealed to the seller. Their setup therefore entails only a binary choice for the consumer. Additionally, the privacy comes at an exogenously given cost, while in our setting costs of privacy are endogenous in that they reduce the fit of the product.

Our paper is also closely related to the literature on agency theory. While classical principal-agent models assume private information to be exogenously given, more recent papers account for the fact that agents often first need to gather information themselves (see, e.g., [Crémer and Khalil, 1992, 1994](#); [Crémer et al., 1998](#); [Kessler, 1998](#)). The paper of [Hoppe and Schmitz \(2013\)](#) provides an experimental test on a fundamental trade-off in agency theory: The principal needs to weigh rent extraction versus an efficient outcome. Additionally, the agent’s information search again has an exogenous cost. In our setup, the consumer chooses the information structure at a solely indirect cost (through a potential price increase), but the seller similarly can promote more efficient trade by giving up rent and setting relatively low prices for higher filter choices.

As already discussed by [Hoppe and Schmitz \(2013\)](#), such setups can also be interpreted as ultimatum games ([Güth et al., 1982](#)). In the classical case, both sides – proposer and responder – know the size of the pie that can be distributed. This setup was modified in subsequent papers such that the pie size is uncertain and only the proposer learns the outcome ([Mitzkewitz and Nagel, 1993](#); [Straub and Murnighan, 1995](#); [Croson, 1996](#); [Güth et al., 1996](#); [Rapoport and Sundali, 1996](#); [Rapoport et al., 1996](#); [Güth and Huck, 1997](#); [Güth and Van Damme, 1998](#); [Huck, 1999](#)). [Kagel et al. \(1996\)](#) and [Harstad and Nagel \(2004\)](#) also consider a case of an informed responder and an uninformed proposer. However, in all these papers, the distribution of the pie size is exogenous. The buyer in our setup can be interpreted as a responder who chooses this distribution actively. The higher the number of filters he chooses, the larger will be the average value of the product offered to him that can than be “split up” between him and the seller. The seller (or proposer) learns the distribution, but the value drawn from the distribution remains

private information of the buyer. Setting a price corresponds to an offer for splitting up the pie and the buyer then can either accept by buying or not. However, our experimental results point to the conclusion that our setup was not perceived as an ultimatum game by the participants in the laboratory, since the propensity of rejected product offers was comparatively low in contrast to usual ultimatum game experiments.

Also related to our setting are experiments that investigate underinvestment due to the hold-up problem. In [Ellingsen and Johannesson \(2004\)](#), for example, one of two players can choose whether to invest unilaterally and in a following bargaining stage, the parties determine the distribution of the outcome generated through this investment. Our setup also constitutes a hold-up problem, but our buyers have a more gradual and indirect way of ‘investing’ through the usage of different filter amounts and the seller has the full bargaining power. In principle, the buyer could use all filters to ensure that he gets offered only the highest valued product, but then the seller could also charge the highest price and capture all the surplus.

Furthermore, there is a related stream of literature on bargaining models with strategic information transmission started by [Crawford and Sobel \(1982\)](#). A recent theoretical working paper in this area ([Hidir and Vellodi, 2018](#)) addresses similar situations as in our project using cheap-talk messages. In contrast to these models, our filter choice is always informative for the seller. The buyer can hide, but not misrepresent, private information.

2 Theoretical Framework

There is one seller (she) and one buyer (he). The seller possesses one good with value normalized to 0 for her, which she wants to sell to the buyer. The buyer’s value v of the product is determined by a value distribution f with support on $[0, 1]$. This is equivalent to assuming that the seller has an (infinitely) large portfolio of products with different valuations for the buyer and the buyer receives one random product.³ With $F(v)$ denoting the cumulative distribution function of v , we assume that the hazard rate $\frac{f(v)}{1-F(v)}$ is non-decreasing and thus f is regular. The buyer’s outside option of not buying the product is also normalized to 0.⁴

The seller provides a search technology for the buyer which we will call *filter*. The filter choice is denoted by $a \in [0, 1]$. By choosing a , the buyer restricts the value function from below,

³Note that there is no objective ex-ante ranking of the products. Differences in buyer valuation stem from taste, not from quality differences.

⁴We assume risk neutrality, but introducing risk aversion on the buyer-side would not change equilibrium predictions.

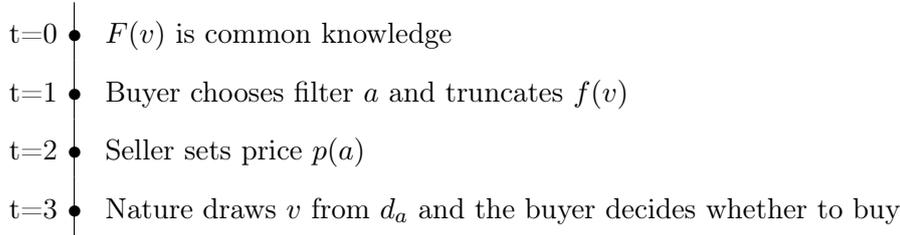
securing him a value of at least a .⁵

The truncated distribution with support on $[a, 1]$ is:

$$d_a(v) = \frac{g(v)}{1 - F(a)}, \quad (1)$$

where $g(v) = f(v)$ for $v \in (a, 1]$ and 0 otherwise.

Crucially, the buyer's choice of a and thus d_a is observed by the seller.⁶ The seller then makes a take-it-or-leave-it offer to the buyer by setting a price p . Finally, nature draws the buyer's valuation for the product v from d_a and the buyer decides on whether to buy or not. The timing of the game is therefore:



We state the following two propositions, which we then test in a lab experiment. The propositions hold for any regular value distribution function $f(v)$ with a non-decreasing hazard rate.

Proposition 1 (Increasing Price):

The optimal seller price is weakly increasing in the filter choice a of the consumer.

Proposition 2 (Strategic Inattention):

The buyer is strategically inattentive and restricts his search optimally such that $a < 1$.

The propositions follow from the subgame-perfect Nash equilibrium of the game, which we derive by using backward induction. The full analysis is provided in Appendix A.1.

Nature draw and buying decision (t=3). Nature draws the buyer's valuation from d_a and the buyer makes his buying decision. We assume that when the buyer is indifferent between

⁵Alternatively this could be modeled as a buyer having a taste parameter θ on a Salop circle with a valuation of $v_{max} - t(\theta - y)$. Here y would be the chosen product and $t(\Delta)$ a function that is non-decreasing in Δ . The corresponding filter choice for the buyer would be to decide on the maximal distance between his target product and the offered product.

⁶If the seller could not observe his filter choice, it would be optimal for the buyer to set $a = 1$.

buying and not buying he will always buy. That is, he will buy as long as $v \geq p$.⁷

Optimal price (t=2). The seller sets a price p given the buyer's choice of a , to maximize his payoff:

$$\Pi_S = p \text{Prob}(v \geq p). \quad (2)$$

The optimal price function for the seller, dependent on a , is

$$p^*(a) = \max\{\hat{p}, a\}, \quad (3)$$

where $\hat{p} \in (0, 1)$ is the optimal price in the unrestricted case ($a = 0$) and thus equal to the standard monopoly price without a search technology available. The price is flat and equal to \hat{p} for $a \leq \hat{p}$ and increasing for $a > \hat{p}$. This proves Proposition 1.

Optimal filter choice (t=1). For the buyer we finally need to find the optimal filter choice, given the seller chooses the optimal price in t=2. The optimal filter choice is

$$a^* = \hat{p}. \quad (4)$$

Since $\hat{p} < 1$, in equilibrium $a^* < 1$, that is the buyer is (partially) inattentive, proving Proposition 2. Intuitively, a lower filter choice than a^* would not lead to a lower price but only to a lower average valuation. On the other hand, the marginal gain from a further increase in filters would be extracted by the seller.

The derived equilibrium filter choice a^* is not efficient and constitutes a hold-up problem in that the buyer restricts his search ($a < 1$), whereas searching fully (choosing $a = 1$) would maximize efficiency.

2.1 Examples

To visualize the above results, we provide two examples with different distributions. In the first one we use a uniform distribution and in the second one a discrete distribution which we also used for the parametrization of the lab experiment.

⁷Relaxing this assumption would not change the results qualitatively.

⁸This price function translates to the optimal price for 'previous customers' in the two-period model discussed in [Fudenberg and Villas-Boas \(2006\)](#).

Example 1. Let v be uniformly distributed with $f(v) = \mathcal{U}(0, 1)$. We obtain $p^* = \max\{\frac{1}{2}, a\}$ and $a^* = \frac{1}{2}$ (see Figure 1).

Example 2. For the lab experiment, we chose a discrete distribution to make it easier to understand for subjects. The above results translate directly to the Binomial distribution, where the filter choice a becomes discrete. We set $f(v) = B(5, 0.5)$. Here, $p^* = \max\{2, a\}$ and $a^* = 2$ (see Figure 2).

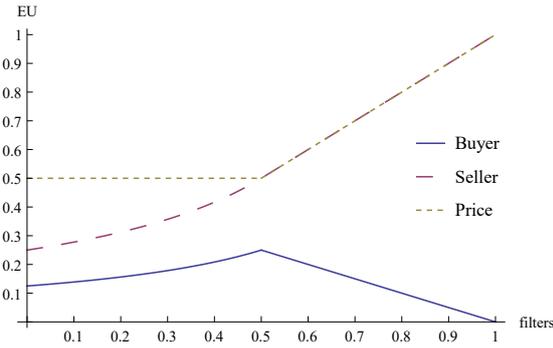


Figure 1: Expected payoffs given filters $a \in [0, 1]$ and $v \sim \mathcal{U}(0, 1)$

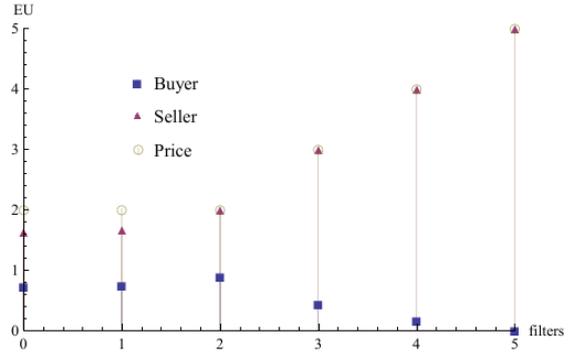


Figure 2: Expected payoffs given filters $a \in \{0, 1, \dots, 5\}$ and $v \sim B(5, 0.5)$

3 Experimental Design

In the experiment, subjects played 15 periods of the game described in Section 2 with the parametrization explained in Example 2 above.⁹ The distribution was presented to the participants as a set of fictitious products with 5 different binary characteristics occurring with the same probability (corresponding to the Binomial distribution $B(5, 0.5)$). In all periods, this set consisted initially of all 2^5 possible combinations of characteristics, which makes a total of 32 products.

Participants were assigned either to the role of a *buyer* or a *seller* and kept this role for the whole experiment. Buyers and sellers were told that in each period the seller can sell at most one good and the buyer has the possibility to buy at most one good, and that the outside option for both buyer and seller is 0.

The steps of the game that the participants play has the following time-line that closely mirrors the timing in the theoretical model:

⁹The full instructions presented to the participants can be found in Appendix A.4.

- t=0 • Buyer receives target product (defining his ideal product characteristics)
- t=1 • Buyer chooses filter a and reduces set of remaining products
- t=2 • Seller sets price dependent on filter choice
- t=3 • One random product of remaining set is offered to buyer
and buyer decides whether to buy

At the beginning of a period, the computer randomly draws from the set of all products to identify the so called *target product* of the buyer. This is the ideal product for the buyer and only he observes the actual characteristics of the target product. The value of each product in the set is determined by the number of matching characteristics with this ideal product. Hence, a product is of value x if x of the 5 binary characteristics are identical with the target product.

The buyer can then set a discrete number of filters $a \in \{0, 1, \dots, 5\}$ to reduce the number of products remaining in the set such that all products left have a minimum value of a for him. The seller then sets a price conditional on the buyer's filter choice. This means that all remaining products will cost the same price. We use the strategy method to elicit the whole price vector of the seller. Hence she sets 6 independent prices, one for each possible filter choice of the buyer. Prices have to be between 0 and the maximum possible valuation of 5 Euros.¹⁰ The strategy method ensures that we obtain pricing data for all possible choices of buyers even when some filter options are only chosen rarely.

Finally, the seller offers one random product of the remaining products from the, potentially reduced, set. Importantly, the seller does not choose the specific product but it is chosen by the computer. The buyer gets a take-it-or-leave-it offer to buy this randomly drawn product for the price that the seller set in the step before. The buyer is free to buy the product or not. When he buys, his payoff is calculated as the value of the product minus the price and the seller gets the price she demanded. When the buyer does not buy, both get 0. Both, seller and buyer, get feedback on the number of filters chosen by the buyer, the actual price set by the seller, the buying decision as well as their own payoffs for this period. The buyer does not obtain information about the rest of the pricing strategy of the seller and the seller does not obtain information about the valuation of the chosen product for the buyer.

For each session we used one of two matching protocols to determine which seller interacts with which buyer over different periods. Under *stranger-matching*, buyers were randomly and

¹⁰All values and prices in the experiment are directly calculated in Euro. Prices can be adjusted in steps of 1 Eurocent.

anonymously re-matched to sellers in the beginning of each period.¹¹ The alternative was a *partner-matching* protocol, where buyers and sellers interacted repeatedly for the whole experiment. We used both protocols to be able to account for different market situations. This way we can also control for the impact of reputation effects, which are only possible under a partner-matching.¹²

Computer task. As a robustness check all subjects participated afterwards in an additional part where they played the same game, but now against the computer. Subjects knew that there was a second part but learned about the details only at the beginning of the computer task where they were informed that they are now playing against a computerized counterpart. As in the main experiment, subjects played for 15 periods and also kept their role as buyers or sellers. Sellers from the main experiment played now against a computer, which choose an amount of filter and made a buying decision. And subjects in the role of a buyer in the main experiment played now against a computer setting a price.

The additional part enables us to disentangle the strategies of buyers and sellers and to observe one-sided behavior. The optimal reaction to a counterpart's non-equilibrium decision might be different from equilibrium choices. In the main experiment, a buyer might not choose the optimal filter choice against a human seller, because he believes that the seller is not sophisticated enough to extract rents. At the same time, a seller might choose prices lower than in equilibrium, because he anticipates that buyers might not buy the product if the net value is too low. Playing against the computer though fixes beliefs about the sophistication of the other party and also controls for possible effects of social preferences.

Sellers set again the full price vector and buyers choose an amount of filters and whether they buy the product in the end. Sellers are told that the computer chooses every filter with equal probability and that it buys when the resulting net value is at least 0. Here the seller should set the optimal price p^* . The buyers from the main experiment now play against a computerized seller that plays the equilibrium strategy with a small random error term attached to the price. They know they play against an algorithm programmed to maximize profits from selling. They are also informed that the algorithm can use their filter decision for the calculation of the prices, but cannot recover information from former periods. Buyers in this part should choose the

¹¹We used matching groups of 8 subjects (4 sellers and 4 buyers).

¹²Reputation effects are considered important in the IO literature. See for example [Kreps and Wilson \(1982\)](#) for a seminal paper.

optimal number of filters a^* .

Procedures. We conducted eight experimental sessions in the BonnEconLab at the University of Bonn in June 2017. Sessions lasted about 75 minutes and we had 192 participants in total (63% female, on average 24.4 years old). The experiment was computerized with oTree (Chen et al., 2016) and participants were invited using hroot (Bock et al., 2014) from the participant pool of the BonnEconLab¹³. For each of the two parts of the experiment (main experiment and computer task), one of the 15 periods was randomly chosen by the computer for payment.¹⁴ The payoffs of both selected periods were paid out in cash to the participants in private directly after the experiment ended. Average earnings summed up to 8.25 €, including a show-up fee of 4 € and payments of post-tests. We collected data on two incentivized post-tests, the “bomb” risk elicitation task (Crosetto and Filippin, 2013; Holzmeister and Pfurtscheller, 2016) as a measure of risk-aversion and the SVO slider measure (Murphy et al., 2011). The experiment ended with a standard post-experimental questionnaire containing demographics.

3.1 Hypotheses

Our hypotheses are based on the two propositions drawn from our model described in Section 2. First we test, in line with Proposition 1, whether the sellers’ price vectors are upward sloping in the filter choice.

Hypothesis 1 *Sellers set higher prices for a higher filter choice.*

Our second main hypothesis is directly derived from Proposition 2, where we test whether buyers choose an optimal filter amount.

Hypothesis 2 *Buyers restrict their search optimally.*

The computer task allows us to see whether facing a human counterpart fundamentally changes behavior compared to the interaction with an automatized computer algorithm. We can observe one-sided behavior and control for social preferences as well as (potentially distorted) beliefs about the sophistication of the other party. We expect that facing a computer should move subject’s behavior closer to the equilibrium predicted by our model.

¹³The only exclusion criterium was more than one no-show in previous experiments.

¹⁴This includes also those periods where the buyer decided against buying the product and both players received the outside option of 0. Hence the computer draw could result in a payment of 0 € for the part.

We also look at the buyer’s final buying decision. Our model assumes implicitly that buyers make rational buying choices and buy whenever their net value is weakly larger than 0. However, it might be that buyers choose not to buy even when they would profit from the deal. A potential reason might be fairness considerations when they find the seller’s price unreasonably high and would make only a small gain. We therefore suspect deviations from the rational choice mainly to happen for small positive net values.

Finally we test whether the matching protocol and potential reputation effects play a role in subjects’ decisions. As explained at the end of Section 2, the equilibrium predicted by the model is not efficient. The possibility to build up a reputation is known as a simplifying factor for market transactions. Fairness considerations with regard to unreasonably high prices might be more apparent under a partner-matching and a seller’s reputation for setting moderate prices might create room for higher filter choices of the buyers and reduce irrational offer rejections. We therefore expect that the partner-matching increases the outcome efficiency compared to the stranger-matching.

4 Results

In this section, we provide the analysis for testing our hypotheses. First note that we did not find any robust qualitative differences between the partner- and the stranger-matching protocol. To provide a concise analysis we thus pool the data from both matching protocols (see Appendix A.2 for separate results per matching protocol).¹⁵ For a discussion of potential reasons why the matching did not significantly influence our results, please refer to Section 6.

4.1 Seller

Figure 3 shows the average price vector in comparison to the price vector predicted by theory.¹⁶ The graph shows that price vectors are on average increasing in the filter choice. This is also supported by the regression depicted in Table 1.¹⁷ This confirms our first hypothesis and shows that sellers use the information about the filter choice to price discriminate.

Result 1 *Prices are increasing in the filter choice.*

¹⁵If not mentioned otherwise, we use matching group averages as individual observations for the sessions with the stranger-matching protocol.

¹⁶For theoretical predictions compare also Figure 2 in Section 2.

¹⁷The analysis still holds when we only look at the actual prices buyers faced in the experiment instead of the whole price vector. The table also shows that the matching protocol has no influence on the price vector.

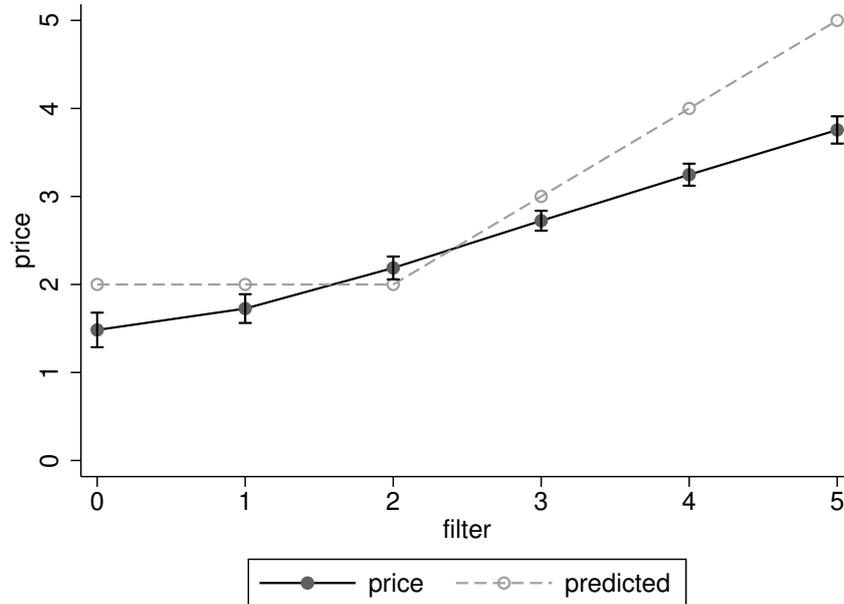


Figure 3: Average price vector (price conditional on filter choice). Bars show 95%-confidence intervals based on subject averages.

Table 1: Regression price vector

	(1)	(2)	(3)
Filter	0.470*** (84.73)	0.470*** (84.73)	0.470*** (84.44)
PARTNER		-0.195 (-1.59)	-0.171 (-1.34)
Controls			YES
Constant	1.325*** (20.96)	1.442*** (15.17)	1.415** (2.72)
Observations	8640	8640	8550

z statistics in parentheses

⁺ $p < .10$, * $p < .05$, ** $p < .01$, *** $p < .001$

Note Mixed effects linear regression with errors nested in matching groups, nested in individuals. Dependent variable is the conditional price according to the strategy method. ‘Partner’ is a dummy taking the value 1 under partner-matching and ‘Filter’ is the filter choice or potential filter choice. Controls contain the period of play, the matching protocol, age, gender, risk-taking, svo-angle, and answers to questions about online shopping from the questionnaire.

Prices are lower than predicted for very low and very high filter choices, but very close to the predicted prices for intermediate filter choices. Thus, although our first hypothesis is confirmed, the average price vector for filter choices above 2 is significantly less steep than the equilibrium prediction (two-sided ttest¹⁸ against predicted slope of 1, $p < 0.001$). This might influence the buyers' optimal filter choices and therefore needs to be taken into account for the analysis of the buyer behavior.

4.2 Buyer

Buyers restrict their search (and choose less than 5 filters) in 75.5% of cases. The average filter amount is 2.69 while the equilibrium prediction is at 2 filters. Therefore buyers choose on average more filters than predicted (two-sided ttest, $p < 0.001$ ¹⁹) and we can only partly confirm our second hypothesis:

Result 2 *Buyers restrict their search behavior, but deviate from the equilibrium prediction.*

This does not necessarily mean that buyers act non-optimally given the situation in the experiment. Given actual price vectors of the sellers differed from the equilibrium price vector, a non-equilibrium filter choice might be the optimal response and lead to higher payoffs for the buyer. Taking the observed average price vector as a benchmark, the optimal filter choice turns out to be 5 filters as can be seen in Figure 4.

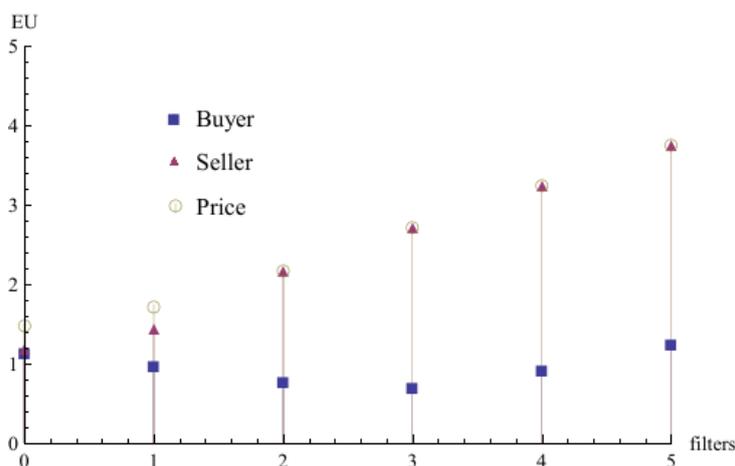


Figure 4: Expected payoffs on filter choice based on observed average price vector

Looking at the histogram of all filter choices in all periods (Figure 5) reveals that buyer choices do not clearly coincide with either the predicted filter choice nor with the optimal filter

¹⁸Average slopes of each individual seller are taken as one observation.

¹⁹Here we treat subject averages as observations.

choice based on the observed average price vector, but are quite diverse. Indeed subjects chose 5 filters most often (24.51%) but the other filter amounts are also chosen between 12% and 17% of the time. The equilibrium prediction of 2 filters was selected in 15 % of all cases.²⁰

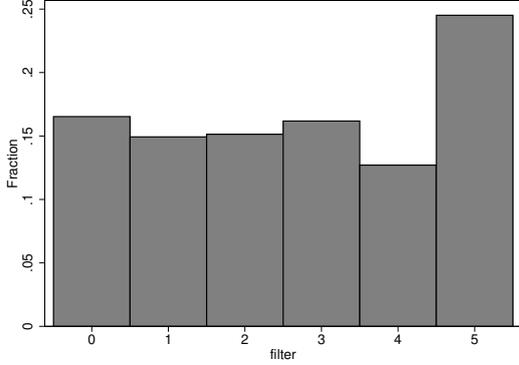


Figure 5: Histogram of all filter choices of all subjects in the main task

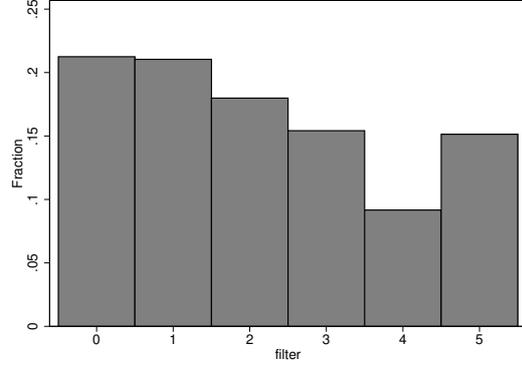


Figure 6: Histogram of all filter choices in the computer task

The same histogram for filter choices in the computer task (Figure 6) suggests that this diversity in behavior is not driven by the non-equilibrium behavior of sellers. Against the computer we can observe one-sided buyer behavior whereby seller behavior is fixed and prices are identical to the equilibrium price vector. Overall, the average filter choice in the computer treatment is with 2.16 lower than in the main treatment and closer to the prediction. However, the predicted number of 2 filters were selected in only 17% of all cases and diversity in choices remains large.

Result 3 *Filter choices are diverse, in the main task as well as in the computer task.*

Comparing Figure 5 and 6, the distributions differ significantly (Kolmogorov–Smirnov test, $p < 0.001$ ²¹) and filter choices move in total in the right direction. There is a shift away from 5 filters (24.5% of choices in the main task, 15.1% in the computer task) and towards choices between 0 and 2 filters (46.6% main, 60.3% computer task). In the computer task, all filters in this range provide a decent expected payoff compared to higher filter choices.²² Thus, buyers restrict their search but don’t search as predicted by the equilibrium.

Altogether behavior in the main as well as in the computer task suggests that buyers do not fully appreciate the impact of the use of the search technology on their payoffs.

²⁰When considering the *stranger*- and *partner*-matching separately we obtain very similar results (see Figure 10 in Appendix A.2).

²¹This calculation is based on all choices. Considering subject averages of filter choices, the difference prevails (Kolmogorov–Smirnov test, $p < 0.001$).

²²See again Figure 2 in Section 2.

4.3 Efficiency and Surplus Distribution

It is not obvious from the separate analysis of buyers and sellers in the experiment which consequences non-equilibrium behavior had on the overall market outcome. Given that the situation constitutes a hold-up problem (not restricting search and buying at any price would maximize efficiency) deviations from the equilibrium could even increase market efficiency.

Efficiency. We therefore investigate efficiency of outcomes in the experiment, whereby efficiency is defined as the sum of payoffs of the seller and the respective buyer. In the experiment, average efficiency is at 2.53 € which is significantly lower than the 2.89 € predicted by the equilibrium (two-sided ttest, $p = 0.0046$).

There are two driving factors for efficiency in our setup: Filter choice and buying probability. Higher filter choices always lead to higher expected efficiency as long as the product is bought. In our theoretical model, buyers always buy the offered product as long as the value is at least as high as the price and the net value is 0 or above. It is conceivable, though, that subjects in the experiment will not buy at very low net values, e.g. due to spite when the seller would get a comparatively high payoff.

Indeed, we find that some subjects reject offered products that would give them a low positive net value. Buyers facing a net value lower than 0.5 € buy only in 17% of the cases, facing a net value above 0.5 € the buying probability is 90%. Thus, efficiency is lower than expected in equilibrium even though 5 filters are chosen most often.²³

Surplus distribution. The decision not to buy the product if the net value is low should reduce the difference between seller and buyer payoffs. Nevertheless, the distribution of surplus between buyers and sellers is quite unequal: A seller earns on average 1.47 € which is significantly more than the average of 1.06 € a buyer gets (sign-rank test with a within-match comparison, $p < 0.001$). But while sellers get on average less payoff than predicted (2 €), buyers do earn more than predicted in equilibrium (0.89 €).

But do buyers also profit from the higher expected efficiency due to higher filter choices in the experiment? To analyze who actually benefits from an increase in filter choice, for buyers we compare average payoffs over all periods in relation to their average filter choice, and for sellers we relate payoffs to the average filter choice that they faced over all periods.

²³Imposing theoretical buying behavior, efficiency in the experiment would be at 3.71 €.

As Figure 7 shows, the buyers' average payoffs are not increasing in the filter choice. In fact it is actually the sellers who profit from an increased use of the search technology provided to the buyers. This dependency is also confirmed in a regression of payoffs on the filter choice (Table 2). The coefficient for filter fails to be significant when we include the interaction term while the other coefficients stay significant. This means that sellers earn less for low filter choices, but they are profiting from an increase in the filter choice while buyer's payoffs do not increase.

Result 4 *Higher filter choices lead to higher welfare. But buyers do not profit from using the filter technology. Only the sellers payoff increases with the buyer's filter choice.*

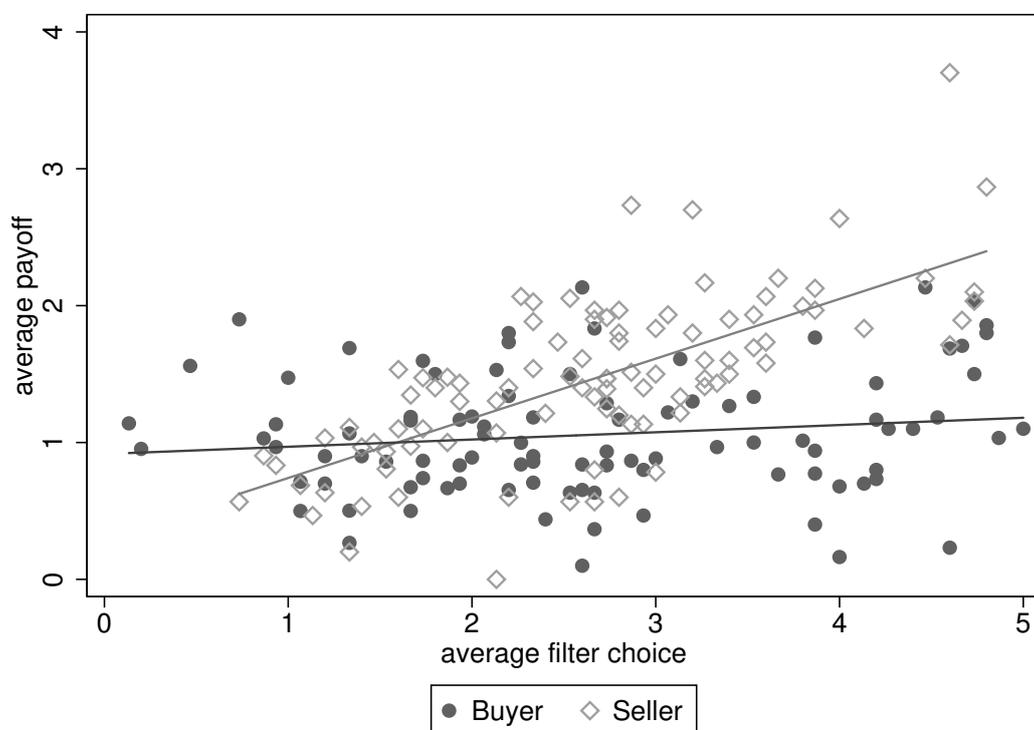


Figure 7: Average payoff on average filter choice (each data point represents one subject and lines are a linear fit)

Table 2: Regression of average payoffs over all periods on average filter choice

	(1)	(2)
Filter	0.199*** (6.67)	0.0529 (1.38)
Seller	0.412*** (5.24)	-0.611*** (-3.53)
Filter \times Seller		0.383*** (5.53)
Constant	0.526*** (5.87)	0.916*** (9.02)
Observations	192	192

z statistics in parentheses

+ $p < .10$, * $p < .05$, ** $p < .01$, *** $p < .001$

Note OLS regression with errors clustered on the matching group level. Dependent variable is average payoffs over all periods. Independent variables are average filter choice / average filter faced, Seller is a dummy for being a seller and Filter \times Seller is an interaction term.

5 Filter Choice Patterns

In this section, we seek to further analyse the observed diversity in filter choices (Result 3). In order to identify potential determinants of this diversity, we will look at patterns of filter choices over time and at how buyers react to prices and payoffs.

As outlined in Section 4.2, the number of filters predicted by equilibrium is chosen about as much as the other options. At the same time, not all buyers keep on constantly varying their filter choice over the 15 periods in the experiment: In 43% of all individual decisions the filter of the last period is chosen again.

First, making use of the panel structure of the experimental data, we analyze determinants for changing the number of filters in general. Table 3 shows that a higher payoff in the previous period makes it significantly less likely to change the number of filters in the current period.²⁴ We also find a significant negative effect of the period of play on the propensity to change the filter. This supports the impression that the filter choices of at least some buyers stabilize over time – although not necessarily at the theoretically optimal number of filters.

Indeed, the individual filter choices over time (see Appendix A.3) reveal quite distinct, but recurring patterns. In the following, we will discuss some individual choice patterns which are each exemplary for a larger group of subjects. This is clearly an ex-post view on the data and

²⁴The results for the computer task can be found in Table 6 in Appendix A.3.

Table 3: Random effects logit regression for determinants of filter change

	Filter change	
Payoff (t-1)	-0.0595* (-2.39)	-0.0573* (-2.30)
Value (t-1)	-0.0443* (-2.06)	-0.0456* (-2.12)
Price (t-1)	0.0282 (1.29)	0.0300 (1.37)
Period	-0.0226*** (-8.02)	-0.0227*** (-8.03)
Controls	YES	
Observations	1344	1344

t statistics in parentheses

⁺ $p < .10$, * $p < .05$, ** $p < .01$, *** $p < .001$

Note Random effects logit regression, marginal effects. Dependent variable is propensity to change filter amount. Independent variables are experienced price and value as well as the payoff in the previous period, ‘Period’ is the period of play. Controls contain the matching protocol, age, gender, risk-taking, svo-angle, and answers to questions about online shopping from the questionnaire.

thus we refrain from making strong claims of discovering specific types of interactions. However, the examples are instructive in showing that the diversity in choices appears within as well as across subjects.

Pattern a (high stable). Both examples in the upper row of Figure 8 reveal an initial phase of varying filter choices for the buyer in the first periods, before the choice of filters stabilizes at the maximum. The prices are such that there is no perfect price discrimination from the seller side. Prices are low enough that they leave some surplus to the buyer. In these examples, participants manage to overcome the hold-up problem and stabilize at an efficient filter choice in the later periods. Interestingly, the left example stems from a condition with partner-matching protocol, while the right pattern appeared under stranger-matching protocol. Under partner-matching, the buyer faces a stable price of a single seller. On the other hand, the buyer in the stranger-matching condition faces varying prices due to the interaction with multiple sellers. Even though prices are varying, his filter choices remain stable, suggesting that varying prices are not a key factor for triggering a change in the filter choice (see insignificant effect of price in Table 3), as long as the share of surplus is high enough for the buyer.

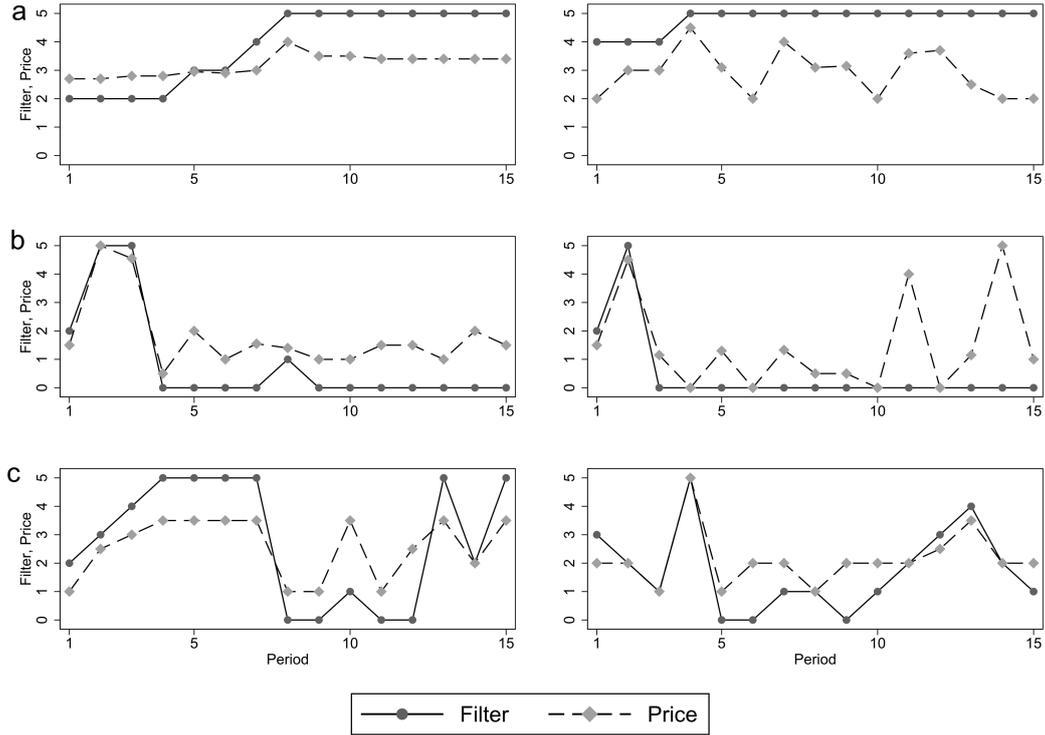


Figure 8: Filter choices for and prices faced by 6 buyers in the experiment. Left column are player ID 218, 404, and 510 from the partner-matching protocol. Right column are ID 118, 622, and 824 from the stranger-matching protocol.

Pattern b (low stable). The examples in the middle row of Figure 8 show a nearly opposite pattern. Buyers choose high filters in the beginning, but also face high prices from the sellers. As a reaction, they nearly completely abandon the use of filters. This leads to a breakdown of the interaction between the seller and the buyer and to low average payoffs for both sides. One interpretation is that, when buyers face strong price discrimination, they consider the use of filters as negative in general and start distrusting the search technology. Therefore, a seller facing a high filter choice is in a critical spot. Extracting rents – setting very high prices for high filter choices – might lead to high payoffs in the short-term. However, this might induce the buyer to reduce his filter choices, which then leads to a loss of efficiency and also a drop in the seller’s payoffs over time. The particular seller might be better off to set intermediate prices and thus give up short-term gains for long-term benefits.

Pattern c (instable). The lower examples in Figure 8 reveal that the observed diversity in choices is not only attributed to differences between subjects – as in the patterns a and b – but that much of the diversity in choices stems from buyers that change their filter choice nearly

every period. They actively explore the strategy space, suggesting that the choice situation is not trivial for buyers. Finding the optimal filter choice is even more challenging in the stranger-matching, where the interaction with varying sellers induces additional fluctuation of prices. Due to frequently changing prices it is not simple to learn the sellers' pricing strategy, or the extent of price discrimination.

6 Discussion

This paper provides a theoretical and experimental analysis of the behavior of sellers and buyers in a situation where the seller can observe the product search behavior of the buyer before setting the price. We show theoretically that prices are increasing in the filter choice and that buyers are strategically inattentive and restrict their search optimally. For the experiment, we apply this setting to a stylized market where buyers can use filters to find better product offers.

Our experimental results are mostly in line with the theory. Sellers price-discriminate by setting higher prices for a higher filter choice, and buyers restrict their search. However, the average price vector is significantly less steep than the equilibrium prediction, and this might influence the buyers' optimal filter choice in the main task. Buyers search more than predicted by the equilibrium, and the distribution of individual filter choices shows that choices are very diverse. This remains true in the computer task where we can observe one-sided behavior. The computer task also confirms that buyers react to the environment in the predicted way, but, while their choices are decent, they are not optimal and also remain diverse.

An exploratory view on individual filter choice patterns over time reveals that this heterogeneous behavior is observable not only across, but also within subjects. This allows insights into potential problems of a buyer-seller interaction in markets where sellers can observe buyers' search behavior, such as online platforms. Some sellers in the experiment seemingly overcome the hold-up problem by charging prices that leave some surplus to the buyer, thus realizing the full potential of search technologies. In other cases, though, the interaction breaks down totally with buyers not choosing filters anymore, probably losing faith in the search technology. This sheds a critical light on the growing interest of the industry in behavioral pricing. While in the short term large gains can potentially be achieved by using consumer data for rent extraction, the long-term consequences are much harder to foresee. A third class of patterns do not show any stabilization of filter choices over time, which suggests that fluctuating prices make the

detection of price-discriminating strategies a challenging task for the buyers.

Our main result, however, is that sellers are profiting from the use of the search technology, while the buyers' average payoffs are not increasing in the filter choice. Additionally, this effect persists in a repeated experimental environment where buyers know about the information transmitted to the seller and where learning would be possible over the periods. This suggests that sellers are able to extract most of the additional surplus generated through the high filter choices of the buyers.

Contrary to our presumption, we find no clear differences between the two matching protocols – stranger and partner – that we used. This means that sellers on average do not gain from the possibility of building a good reputation. While buyers seem not to profit from the use of higher filters, we assume that sellers take into account that buyers can reject the offer in each period. The decision not to buy a product in the end already restricts the sellers' in the one-shot interaction, not leaving much room for improvement due to reputation. The voluntary buying decision might therefore function as an instrument for a natural upper bound to price discrimination.

The example we use in the paper and to set up our experiment is based on the idea that sellers observe the filter choice of buyers. This filter-frame was introduced to capture the theoretical model in a familiar decision context that can be found on actual internet platforms. On the other hand, we restrict the search possibilities to one particular option, while in the field buyers can use many different ways to search for a product. Additionally, sellers in our setup can set prices directly in response to the filter choice. In the field, online platforms might use information about buyers in a more indirect way. We consider our basic setting to be applicable to different types of information transmission, as long as the buyer reveals more information with more extensive product search. Thus, our paper provides a first step for further research, including more complex types of information search.

There are, of course, further differences from our experimental setup to the field. Our observed patterns show that price discrimination might lead to a loss of trust in the search technology. Online platforms have to take additional negative impacts of price discrimination into account, e.g., the risk of total boycotts. Also, buyers in the experiment knew that their search information can be used to price-discriminate. In online markets with fast-moving prices, price discrimination is harder to detect because consumers cannot observe whether information is transmitted and used for price-setting. We therefore assume that our findings serve as an

upper bound for the consumers' ability to be strategically inattentive.

In this paper, we use a monopoly setting and discuss the interaction of a single buyer and a single seller. This allows us to concentrate on the main aspect of information transmission from the buyer to the seller through disclosed search behavior. A natural next step would be to introduce competition (Fudenberg and Tirole, 2000; Chen and Zhang, 2009) between multiple buyers and multiple sellers. In this case, other aspects of optimal search behavior also become relevant. As an example, lock-in effects would be possible, as consumers might only consult a limited amount of platforms before making their buying decision.

The relationship between consumers and sellers is currently changing due to technological developments in online markets. There is a rising awareness for the potential need of consumer protection in connection to these emerging markets. The European Commission adopted the Digital Single Market Strategy already back in 2016. Since the same year, the guidelines for the "Directive on Unfair Commercial Practices" (UCPD)²⁵ of the European Commission contain also guidelines how to apply UCPD to online platforms. Recently, they additionally set up an expert group on the online platform economy to support the Commission in better regulating these markets.

A single study cannot decide whether extended consumer protection on online platforms is indicated. But our paper allows first behavioral insights into how consumers cope with these developments, shedding a critical light on the gains from a multitude of product information that seemingly comes for free. Our results suggest that, without further intervention, it could be the sellers and not the buyers who profit from technological advances in information technology.

²⁵http://ec.europa.eu/justice/consumer-marketing/files/ucp_guidance_en.pdf

References

- Acquisti, A., C. Taylor, and L. Wagman (2016). The economics of privacy. *Journal of Economic Literature* 54(2), 442–92.
- Bergemann, D., B. Brooks, and S. Morris (2015). The limits of price discrimination. *American Economic Review* 105(3), 921–57.
- Bock, O., I. Baetge, and A. Nicklisch (2014). hroot: Hamburg registration and organization online tool. *European Economic Review* 71, 117–120.
- Caplin, A., M. Dean, and J. Leahy (2017). Rationally inattentive behavior: Characterizing and generalizing shannon entropy. Technical report, National Bureau of Economic Research.
- Carrasco, V., V. F. Luz, N. Kos, M. Messner, P. Monteiro, and H. Moreira (2018). Optimal selling mechanisms under moment conditions. *Journal of Economic Theory*.
- Chen, D., M. Schonger, and C. Wickens (2016). otree – an open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance* 9, 88–97.
- Chen, Y. and J. Zhang (2009). Dynamic targeted pricing with strategic consumers. *International Journal of Industrial Organization* 27(1), 43–50.
- Condorelli, D. and B. Szentes (2018). Information design in the hold-up problem. *Journal of Political Economy*, forthcoming.
- Crawford, V. and J. Sobel (1982). Strategic information transmission. *Econometrica: Journal of the Econometric Society*, 1431–1451.
- Crémer, J. and F. Khalil (1992). Gathering information before signing a contract. *The American Economic Review*, 566–578.
- Crémer, J. and F. Khalil (1994). Gathering information before the contract is offered: The case with two states of nature. *European Economic Review* 38(3-4), 675–682.
- Crémer, J., F. Khalil, and J.-C. Rochet (1998). Strategic information gathering before a contract is offered. *Journal of Economic Theory* 81(1), 163–200.
- Crosetto, P. and A. Filippin (2013). The bomb risk elicitation task. *Journal of Risk and Uncertainty* 47(1), 31–65.

- Croson, R. T. (1996). Information in ultimatum games: An experimental study. *Journal of Economic Behavior & Organization* 30(2), 197–212.
- Dengler, S. and J. Prüfer (2018). Consumers’ privacy choices in the era of big data.
- Ellingsen, T. and M. Johannesson (2004). Is there a hold-up problem? *Scandinavian Journal of Economics* 106(3), 475–494.
- Fudenberg, D. and J. Tirole (2000). Customer poaching and brand switching. *RAND Journal of Economics*, 634–657.
- Fudenberg, D. and M. Villas-Boas (2006). Behavior-based price discrimination and customer recognition. *Handbook on economics and information systems* 1, 377–436.
- Güth, W. and S. Huck (1997). From ultimatum bargaining to dictatorship—an experimental study of four games varying in veto power. *Metroeconomica* 48(3), 262–299.
- Güth, W., S. Huck, and P. Ockenfels (1996). Two-level ultimatum bargaining with incomplete information: An experimental study. *The Economic Journal*, 593–604.
- Güth, W., R. Schmittberger, and B. Schwarze (1982). An experimental analysis of ultimatum bargaining. *Journal of Economic Behavior & Organization* 3(4), 367–388.
- Güth, W. and E. Van Damme (1998). Information, strategic behavior, and fairness in ultimatum bargaining: An experimental study. *Journal of mathematical Psychology* 42(2-3), 227–247.
- Hannak, A., G. Soeller, D. Lazer, A. Mislove, and C. Wilson (2014). Measuring price discrimination and steering on e-commerce web sites. In *Proceedings of the 2014 conference on internet measurement conference*, pp. 305–318. ACM.
- Harstad, R. and R. Nagel (2004). Ultimatum games with incomplete information on the side of the prosper: An experimental study. *Uadernos de Economía* 27(75), 37–74.
- Hart, O. and J. Tirole (1988). Contract renegotiation and coasian dynamics. *The Review of Economic Studies* 55(4), 509–540.
- Hidir, S. and N. Vellodi (2018). Personalization, discrimination and information revelation. *Working Paper*.

- Holzmeister, F. and A. Pfurtscheller (2016). otree: The “bomb” risk elicitation task. *Journal of Behavioral and Experimental Finance* 10, 105–108.
- Hoppe, E. I. and P. W. Schmitz (2013). Contracting under incomplete information and social preferences: An experimental study. *Review of Economic Studies* 80(4), 1516–1544.
- Huck, S. (1999). Responder behavior in ultimatum offer games with incomplete information. *Journal of Economic Psychology* 20(2), 183–206.
- Hunold, M., R. Kesler, and U. Laitenberger (2018). Hotel rankings of online travel agents, channel pricing, and consumer protection. Technical report, ZEW Discussion Papers.
- Ichihashi, S. (2019). Online privacy and information disclosure by consumers. *American Economic Review* (forthcoming).
- Kagel, J. H., C. Kim, and D. Moser (1996). Fairness in ultimatum games with asymmetric information and asymmetric payoffs. *Games and Economic Behavior* 13(1), 100–110.
- Kessler, A. S. (1998). The value of ignorance. *The Rand Journal of Economics*, 339–354.
- Kreps, D. M. and R. Wilson (1982). Reputation and imperfect information. *Journal of economic theory* 27(2), 253–279.
- Martin, D. (2016). Rational inattention in games: Experimental evidence. *Working Paper*.
- Martin, D. (2017). Strategic pricing with rational inattention to quality. *Games and Economic Behavior* 104, 131–145.
- Mikians, J., L. Gyarmati, V. Erramilli, and N. Laoutaris (2012). Detecting price and search discrimination on the internet. In *Proceedings of the 11th ACM Workshop on Hot Topics in Networks*, pp. 79–84. acm.
- Mitzkewitz, M. and R. Nagel (1993). Experimental results on ultimatum games with incomplete information. *International Journal of Game Theory* 22(2), 171–198.
- Montes, R., W. Sand-Zantman, and T. Valletti (2018). The value of personal information in online markets with endogenous privacy. *Management Science*.
- Murphy, R., K. Ackermann, and M. Handgraaf (2011). Measuring social value orientation. *Judgment and Decision Making* 6(8), 771–781.

- Rapoport, A. and J. A. Sundali (1996). Ultimatums in two-person bargaining with one-sided uncertainty: Offer games. *International Journal of Game Theory* 25(4), 475–494.
- Rapoport, A., J. A. Sundali, and D. A. Seale (1996). Ultimatums in two-person bargaining with one-sided uncertainty: Demand games. *Journal of Economic Behavior & Organization* 30(2), 173–196.
- Roesler, A.-K. and B. Szentes (2017). Buyer-optimal learning and monopoly pricing. *American Economic Review* 107(7), 2072–2080.
- Sims, C. (2003). Implications of rational inattention. *Journal of monetary Economics* 50(3), 665–690.
- Sims, C. (2006). Rational inattention: Beyond the linear-quadratic case. *The American economic review* 96(2), 158–163.
- Straub, P. G. and J. K. Murnighan (1995). An experimental investigation of ultimatum games: Information, fairness, expectations, and lowest acceptable offers. *Journal of Economic Behavior & Organization* 27(3), 345–364.
- Villas-Boas, M. (2004). Price cycles in markets with customer recognition. *RAND Journal of Economics*, 486–501.
- Vissers, T., N. Nikiforakis, N. Bielova, and W. Joosen (2014). Crying wolf? on the price discrimination of online airline tickets. In *7th Workshop on Hot Topics in Privacy Enhancing Technologies (HotPETs 2014)*.

A Appendix

A.1 Theory – extended

Optimal Price

The seller sets a price dependent on the buyer's choice of a to maximize his payoff

$$\Pi_S = p \text{Prob}(v \geq p). \quad (5)$$

Case 1. Assume $p \leq a$. Clearly, $\text{Prob}(v \geq p) = 1$ and thus the seller increases his profit by setting $p = a$.

Case 2. Assume $p > a$. Then,

$$\Pi_S = p \int_p^1 \frac{g(v)}{1 - F(a)} dv = p \int_p^1 \frac{f(v)}{1 - F(a)} dv = \frac{1 - F(p)}{1 - F(a)} p \quad (6)$$

$$\frac{\partial \Pi_S}{\partial p} = \frac{1 - F(p)}{1 - F(a)} - \frac{F'(p)}{1 - F(a)} p \stackrel{!}{=} 0. \quad (7)$$

This reduces to:

$$p \frac{f(p)}{1 - F(p)} = 1. \quad (8)$$

Because the hazard rate $\frac{f(x)}{1 - F(x)}$ is assumed to be non-decreasing in x , there is a unique solution \hat{p} with $\hat{p} \in (0, 1)$, solving the above equation (8). Taking both cases together, we get $p^*(a) = \max\{\hat{p}, a\}$. This proves Proposition 1.

Optimal Filter Choice

For the buyer, we need to find the optimal filter choice if the seller chooses the optimal price. We again look at the two cases separately.

Case 1. Assume $a \leq \hat{p}$ and thus $p = p^* = \hat{p}$. Then,

$$\Pi_B = E[v - p | v \geq p] = \int_a^1 \frac{f(v)}{1 - F(a)} (v - p) dv \quad (9)$$

$$\frac{\partial \Pi_B}{\partial a} = \frac{f(a)}{(1 - F(a))^2} \int_a^1 f(v)(v - p) dv > 0. \quad (10)$$

Case 2. Assume $a > \hat{p}$. Then, $p = p^* = a$:

$$\Pi_B = \int_a^1 \frac{f(v)}{1 - F(a)} (v - a) dv = \frac{1 - a - \int_a^1 F(v) dv}{1 - F(a)}. \quad (11)$$

Again, given that the hazard rate $\frac{f(x)}{1 - F(x)} < 1$, we see that the first derivative is negative:

$$\frac{\partial \Pi_B}{\partial a} = -1 + \frac{f(a)}{1 - F(a)} \frac{1 - a - \int_a^1 F(v) dv}{1 - F(a)} < 0. \quad (12)$$

Taken together, we get that the optimal choice $a^* = \hat{p} < 1$. That is, in equilibrium the buyer restricts his search, proving Proposition 2.

A.2 Results – extended

Here, the graphs and regressions from the main body of the paper are displayed separately for the stranger-matching and the partner-matching protocol. Also the price vector regression is repeated for the computer task.

Seller

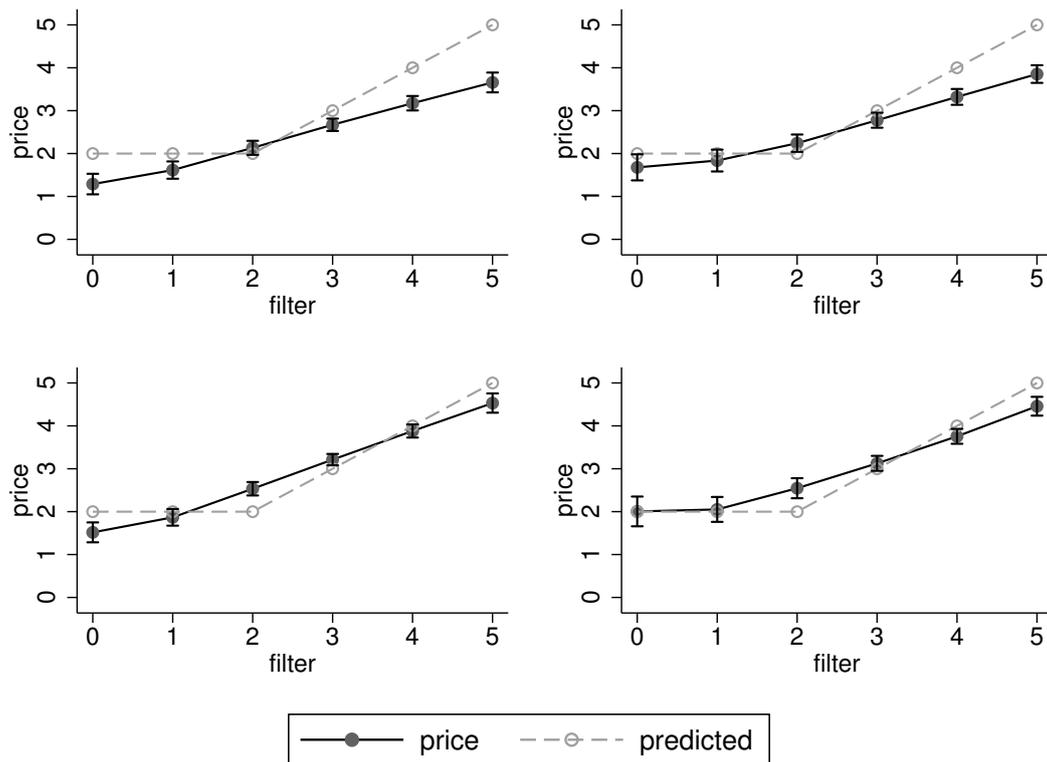


Figure 9: Average price vectors; left: partner-matching, right: stranger-matching; top: main task, bottom: computer task. Bars represent 95%-confidence intervals based on subject averages.

Table 4: Regression price vector (computer task)

	(1)	(2)	(3)
Filter	0.568*** (117.19)	0.568*** (117.19)	0.568*** (116.40)
PARTNER		-0.0667 (-0.59)	-0.114 (-0.94)
Controls			YES
Constant	1.538*** (26.58)	1.571*** (19.45)	0.722 (1.35)
Observations	8640	8640	8550

z statistics in parentheses

+ $p < .10$, * $p < .05$, ** $p < .01$, *** $p < .001$

Note Mixed effects linear regression for the computer task with errors nested in matching groups, nested in individuals. Dependent variable is the conditional price according to the strategy method. ‘Partner’ is a dummy taking the value 1 under partner-matching and ‘Filter’ is the filter choice or potential filter choice. Controls contain the period of play, the matching protocol, age, gender, risk-taking, svo-angle, and answers to questions about online shopping from the questionnaire.

Buyer

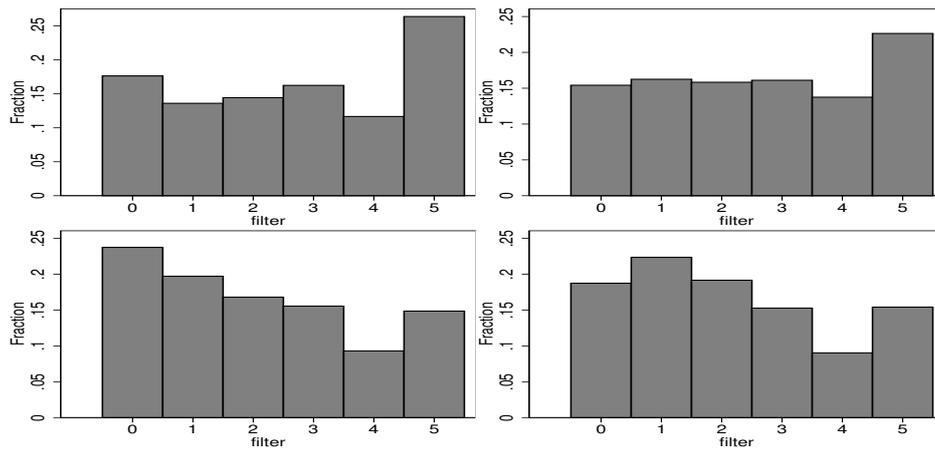


Figure 10: Histograms of filter choices; left: partner-matching, right: stranger-matching; top: main task, bottom: computer task. In bottom histograms, subjects are split up after matching in main task.

Efficiency and Surplus Distribution

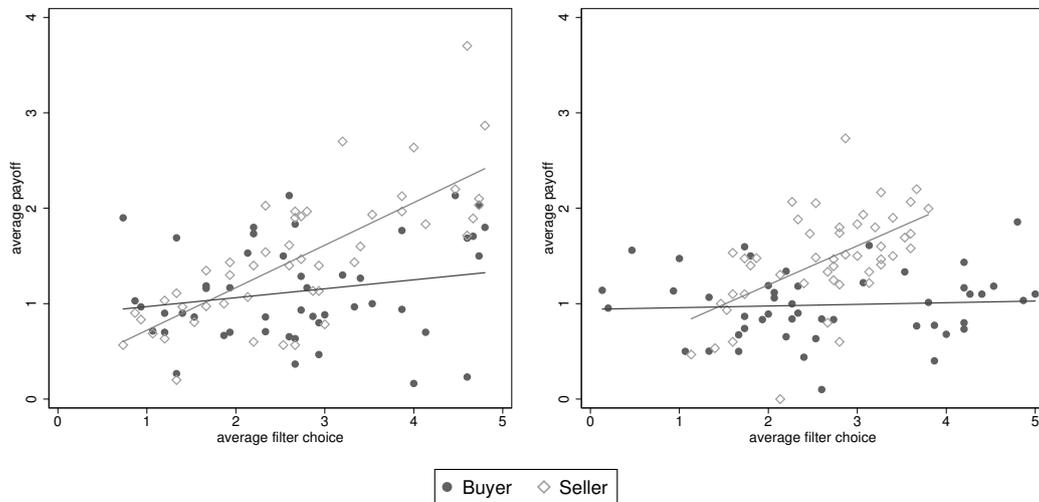


Figure 11: Average payoff on average filter choice; left: partner-matching, right: stranger-matching

Table 5: Regression of average payoffs over all periods on average filter choice

	stranger	partner
Filter	0.0174 (0.51)	0.0936 (1.31)
Seller	-0.563* (-2.25)	-0.600* (-2.33)
Filter \times Seller	0.392*** (4.89)	0.352** (3.22)
Constant	0.941*** (9.77)	0.876*** (4.60)
Observations	96	96

z statistics in parentheses. $^+ p < .10$, $^* p < .05$, $^{**} p < .01$, $^{***} p < .001$

Note OLS regression with errors clustered on the matching group level. Dependent variable is average payoffs over all periods. Independent variables are average filter choice / average filter faced, Seller is a dummy for being a seller and Filter \times Seller is an interaction term.

A.3 Individual Filter Choices

Main Task

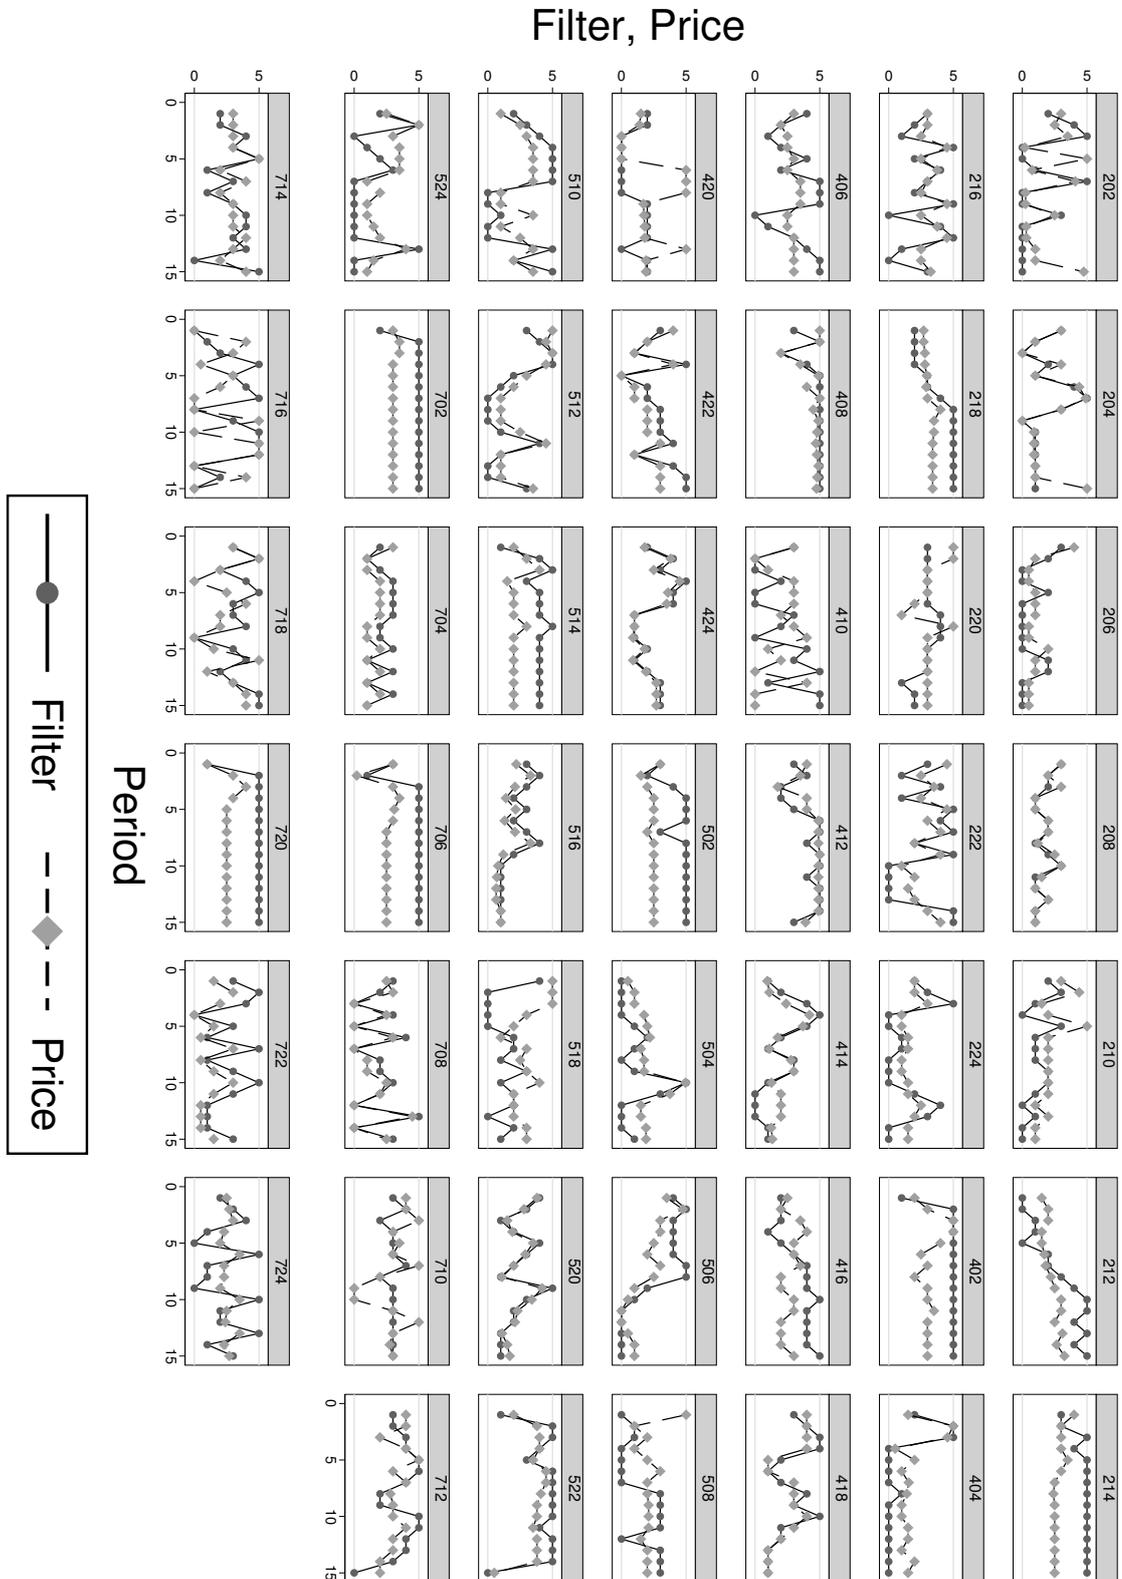


Figure 12: Individual filter choices and realized prices for buyers (partner-matching protocol)

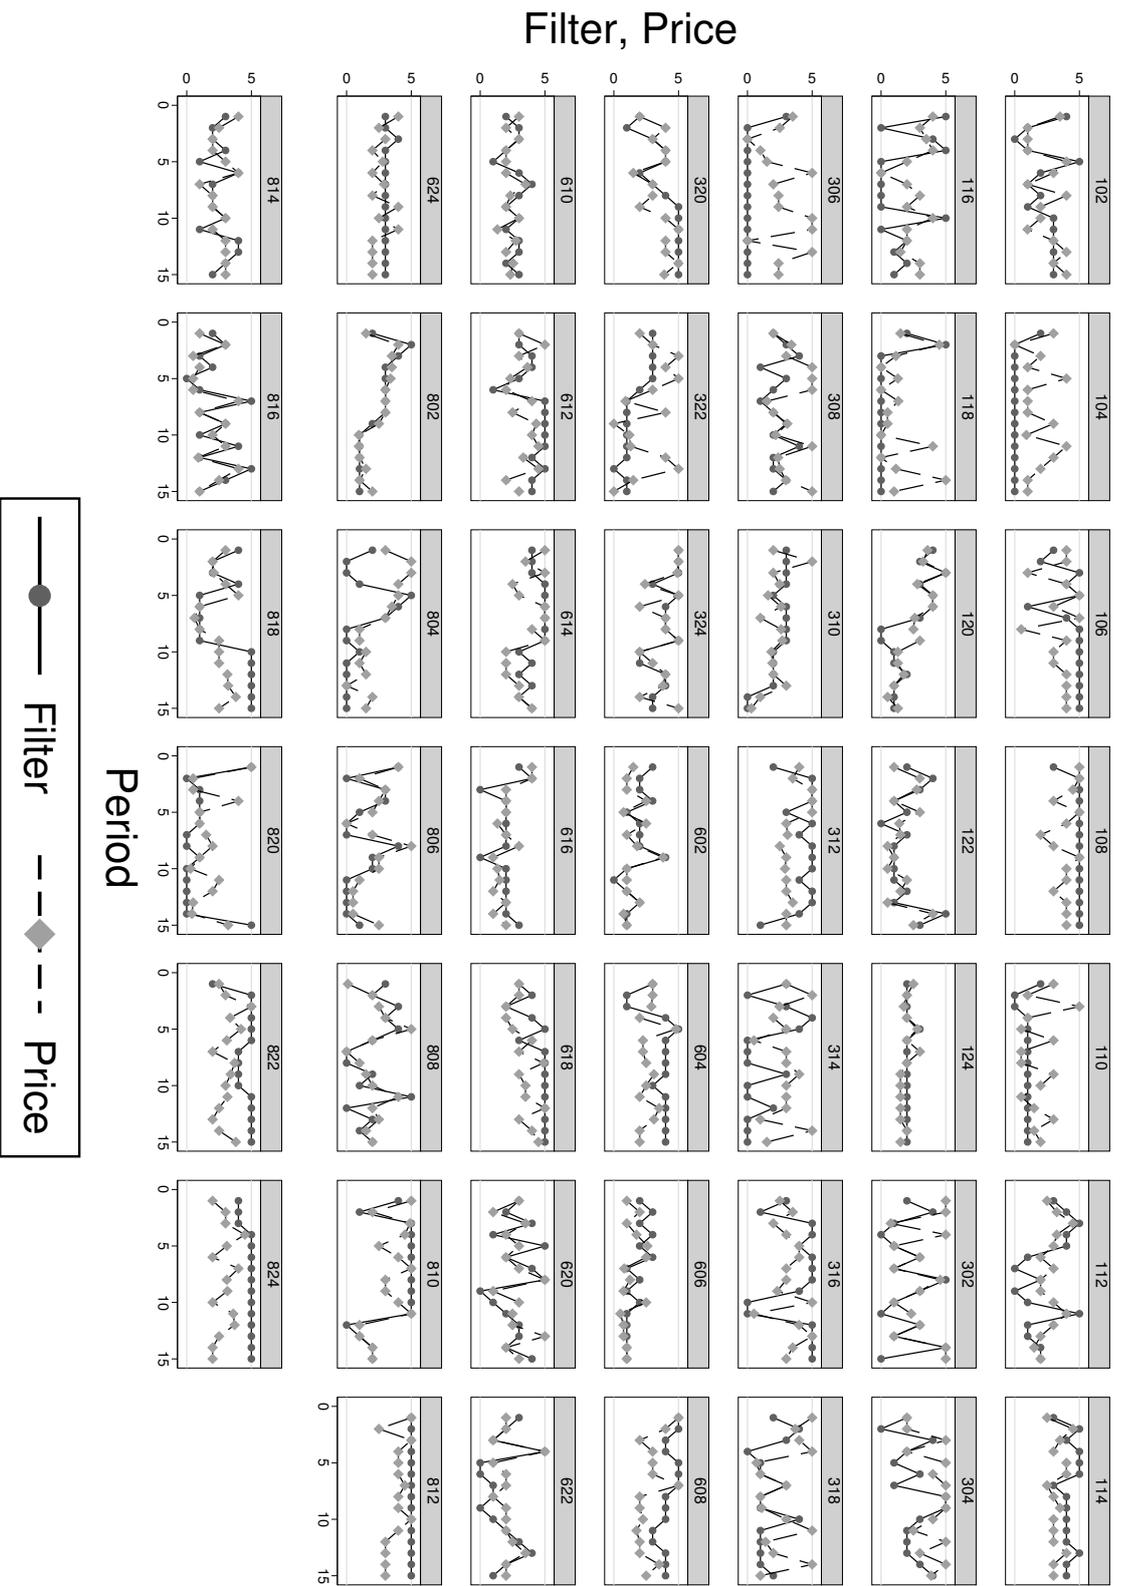


Figure 13: Individual filter choices and realized prices for buyers (stranger-matching protocol)

Computer Task

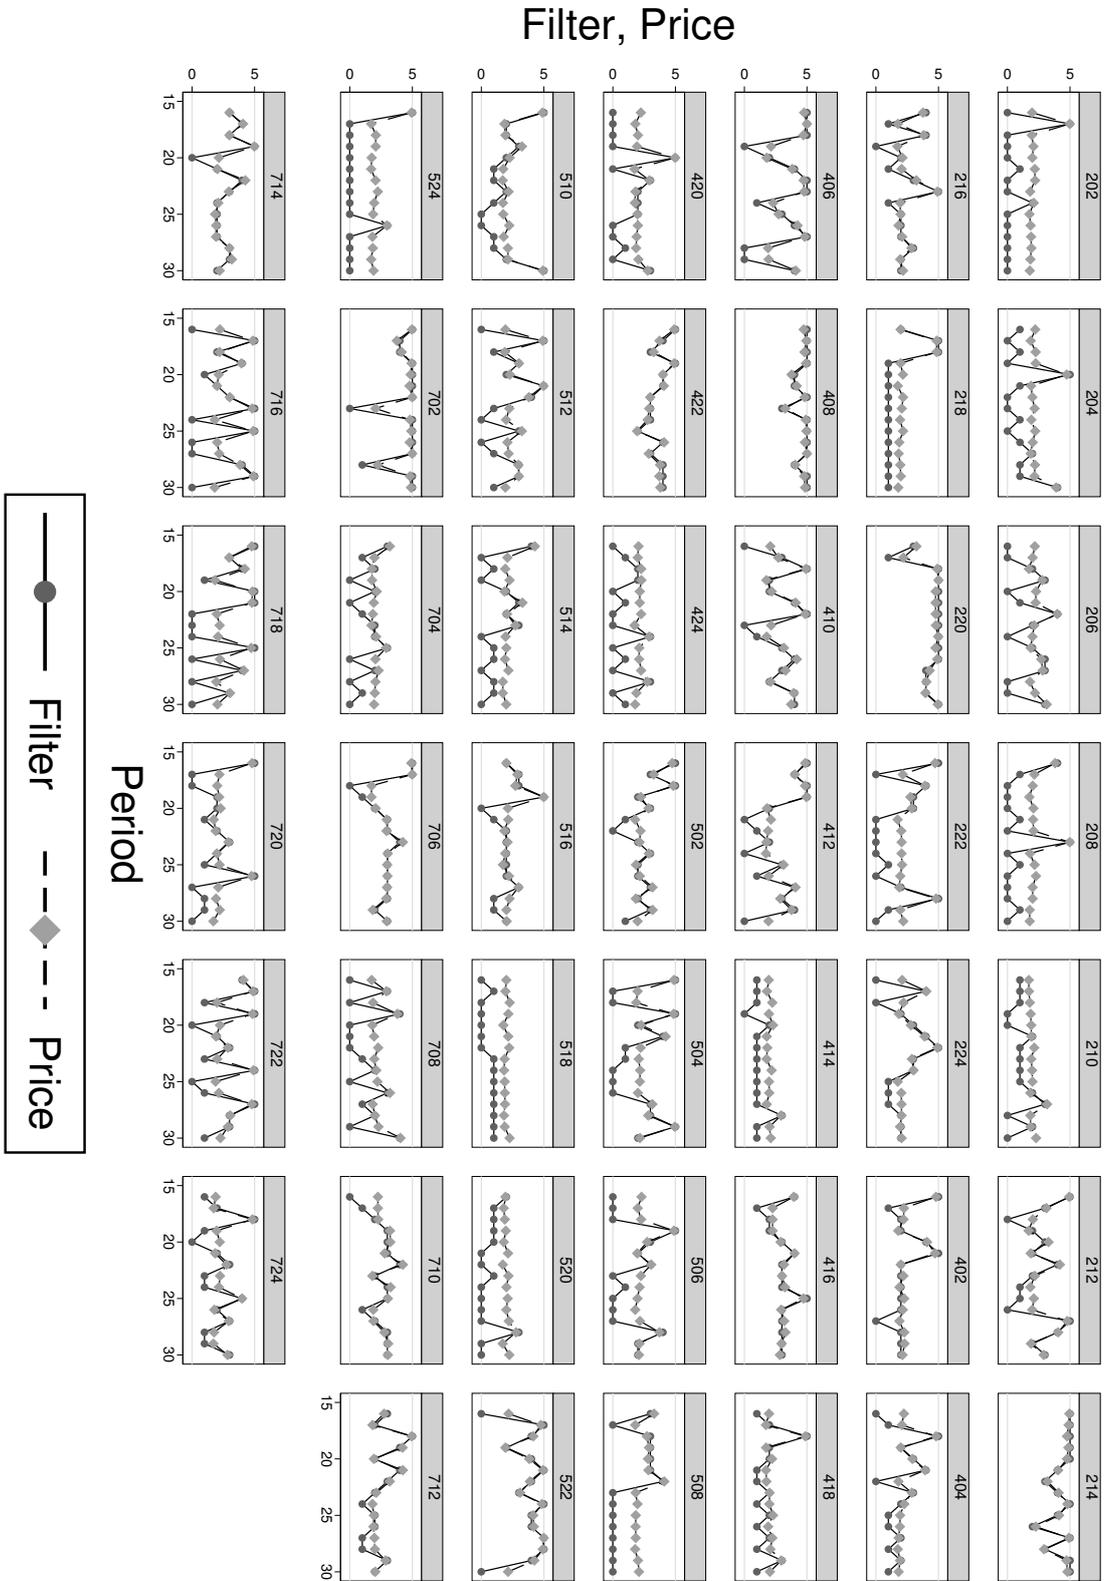


Figure 14: Computer task - Individual filter choices and realized prices for buyers (partner-matching protocol in main task)

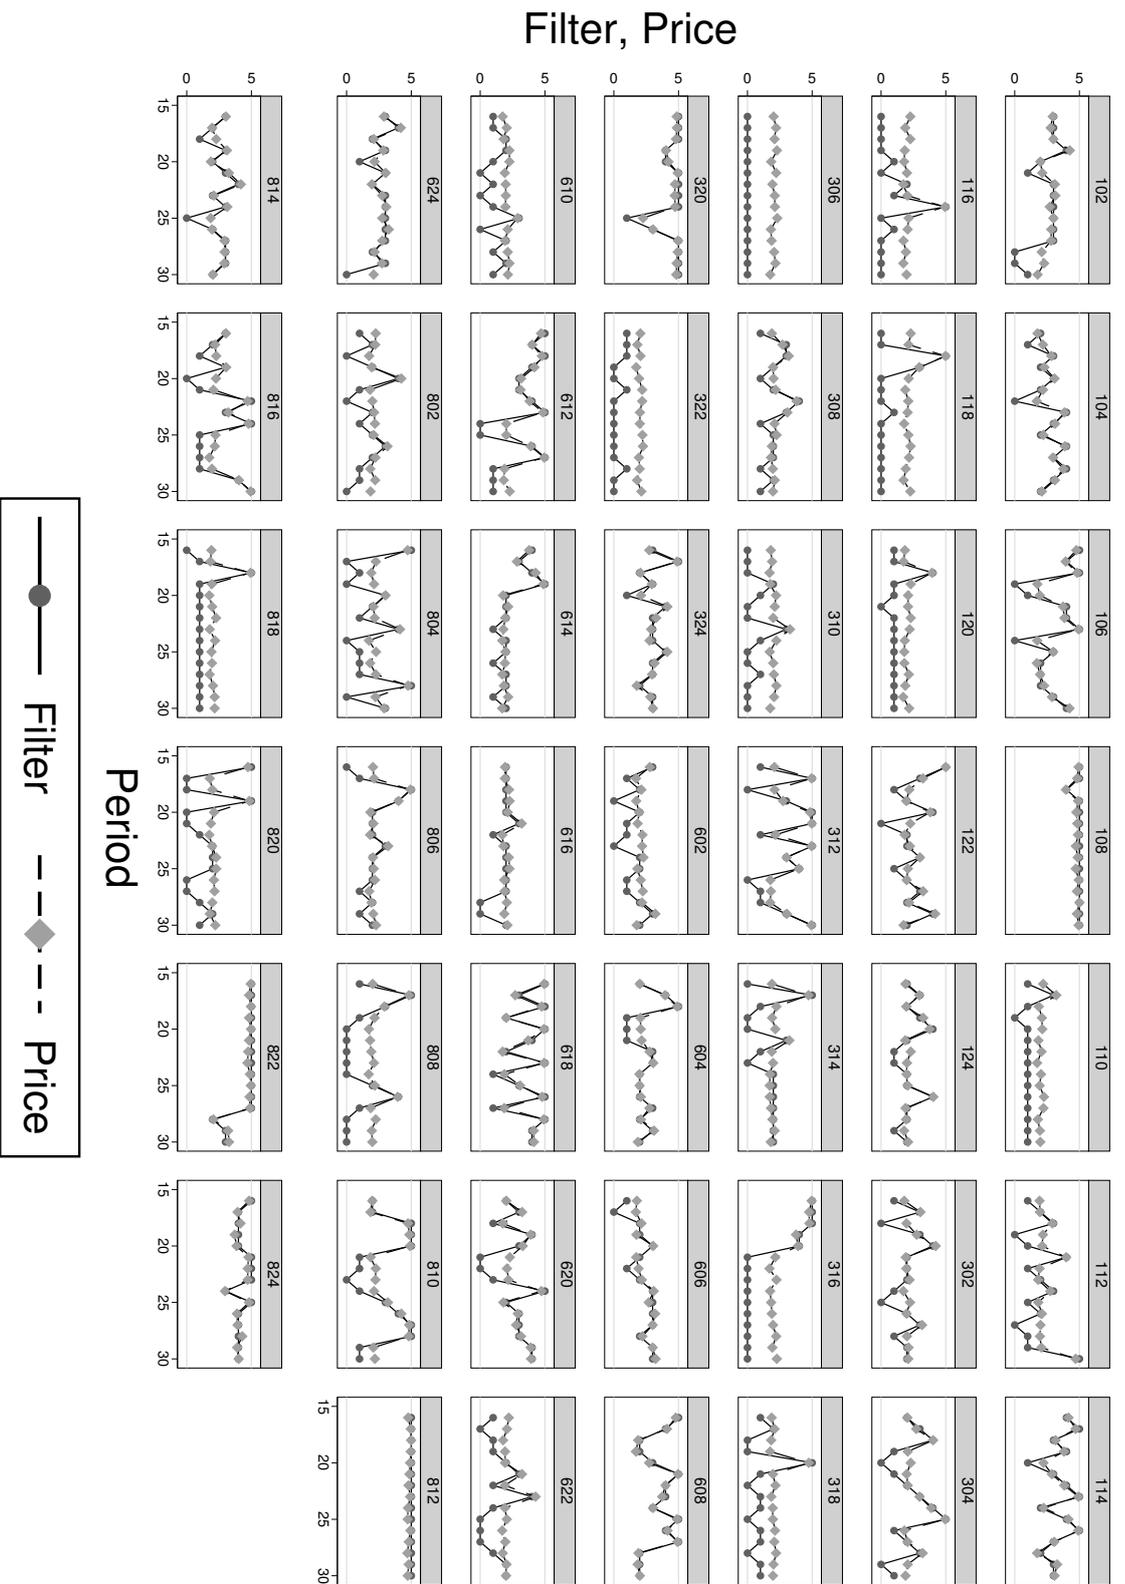


Figure 15: Computer task - Individual filter choices and realized prices for buyers (stranger-matching protocol in main task)

Table 6: Random effects logit regression for determinants of filter change in the computer task.

	(1)	(2)
Payoff (t-1)	-0.0936* (-2.22)	-0.0918* (-2.18)
Value (t-1)	-0.0265 (-0.76)	-0.0281 (-0.80)
Price (t-1)	0.0446 (1.10)	0.0461 (1.14)
Period	-0.0134*** (-4.78)	-0.0133*** (-4.78)
Controls		YES
Observations	1344	1344

t statistics in parentheses

⁺ $p < .10$, * $p < .05$, ** $p < .01$, *** $p < .001$

Note Random effects logit regression, marginal effects. Dependent variable is propensity to change filter amount. Independent variables are experienced price and value as well as the payoff in the previous period, 'Period' is the period of play. Controls contain the matching protocol, age, gender, risk-taking, svo-angle, and answers to questions about online shopping from the questionnaire.

A.4 Instructions

*Note: Instructions displayed here are a translation into English. Original instructions were in German and are available from the authors upon request.*²⁶

Welcome to our experiment!

You are about to take part in an economic experiment that is financed by the Max Planck Society. It is therefore very important that you read the following instructions carefully. You will receive 4 Euro for showing up to this experiment. During the experiment, you will be given the chance to earn further sums of money. The exact amount will depend both on your own decisions and on the decisions made by the other participants in the experiment, as well as on chance. All sums mentioned during the experiment are calculated directly in Euro. After the experiment, you will be asked to fill in a brief questionnaire. Before you leave, all sums of money you have earned will be paid out to you in cash in Euro.

Please stop talking now, switch off your mobile phone, and remove from your desk anything you do not need for this experiment. Disobeying these rules will lead to exclusion from the experiment and from all payments.

Today's experiment consists of two parts. You will receive the instructions to the individual parts just before each respective part begins. The decisions you make in both parts will have no impact on the respective other part or on the payments you can receive in that other part.

In the following paragraphs, we will describe the exact procedure of the first part of this experiment. At the end of this introductory information, we will ask you please to answer some comprehension questions on your computer screen, which are meant to familiarize you with the decision situation.

Should you have any questions, please raise your hand quietly and ask only us. We will then come to you and answer your questions individually.

²⁶We thank Brian Cooper from the MPI for Collective Goods for the translation.

Description of the First Part

This part of the experiment consists of 15 rounds. At the end of the experiment, exactly one round will be chosen randomly by the computer for payoff. Since you will only be told at the end of the experiment which round is payoff-relevant, you should make your decision carefully in each round.

Before the beginning of the first round, the computer will assign you either the role of buyer or that of seller. You will keep this role during the entire first part; it will therefore not change.

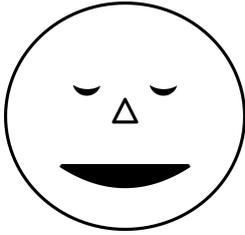
[*Only under stranger-matching protocol:* Further, the computer will rematch, in each new round, one buyer and one seller. In each of the 15 rounds, you will interact with a randomly chosen buyer or seller. At no point, neither during nor after the experiment, will you be told with which other people you have been matched.]

[*Only under partner-matching protocol:* In addition, the computer will randomly match a buyer and a seller at the beginning of the experiment. In all 15 rounds, you will always interact with the same buyer or seller. At no point, neither during nor after the experiment, will you be told with which other people you have been matched.]

In each round, buyers have the chance to buy a product from a large selection of products. The products are depicted through faces that differ in a maximum of 5 characteristics. These characteristics are:

- A white or gray color
- A closed or open left eye
- A closed or open right eye
- A triangular or diamond-shaped nose
- A closed or open mouth.

A total of exactly 32 different products can be made from these characteristics, and each of these possible products occurs just once. For example, a product can be a white face with a



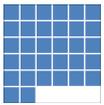
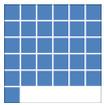
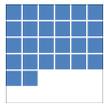
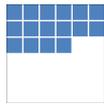
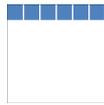
closed left eye, a closed right eye, a triangular nose, and an open mouth. This product would then look as follows:

In each round, the buyer is assigned a target product at random by the computer. At no point during the experiment can the seller see this product. The target product would be the optimal product for the buyer and has a value of 5 € for the buyer. No other product has a equivalently high value. The value of each product is determined from the number of matching characteristics with the target product. If, for example, a product distinguishes itself from the target product only by a different nose, then it has a value of 4 € for the buyer. If, on the other hand, a product differs from the target product in all 5 characteristics, then the value of this product is 0 € for the buyer. The products therefore have a value of exactly 0 €, 1 €, 2 €, 3 €, 4 €, or 5 € for the buyer.

At the beginning of each round, the buyer has the chance to diminish the number of products by choosing a number of filters. The buyer may choose between 0, 1, 2, 3, 4, or 5 filters. The number of filters determines the minimum value of the remaining products. In other words, those products remain that have at least as many matching characteristics with the target product as the number of filters chosen. Or, put differently, products with a lower value for the customer are dropped.

Example: If a buyer chooses 0 filters, then all possible 32 products remain. If, for example, the buyer chooses 4 filters, only those products remain that have a value of 4 € or 5 € for the buyer (i.e., a minimum value of 4 €), and all products with a value of 3 € or less are dropped. The number of products remaining – dependent on the number of filters – and the corresponding minimum value can be examined on the decision screen.

Information on the decision screen:

Filter	0	1	2	3	4	5
Remaining products						
Minimum value	0	1	2	3	4	5

In each round, the seller determines a price between 0 € and 5 € for every possible filter choice of the buyer (0, 1, 2, 3, 4, 5). This price can be chosen quite precisely, up to two positions after the decimal point. The seller is also shown the above information. However, the seller is not told the buyer's target product.

The actual filter choice then determines the price that is chosen. This price is then valid for all other products. Hence, if the buyer wishes to buy a product, this price must be paid. The buyer is told the price that was determined by the seller only after the buyer has opted for a number of filters. The buyer is only told the price corresponding to his or her choice of filters, rather than the prices the seller has chosen for the possible other filters.

After the buyer has been told the price for the remaining products, the computer randomly chooses a product from the set of remaining products. The buyer can then either buy or not buy this selected product. The buyer is shown the selected product by the computer, as well as the payoff that beckons in this round if he or she buys the product. The seller is not told which product has been randomly chosen and hence does not know that product's value either.

At the end of each round, the buyer and the seller are shown how high the payoff for this round is. In addition, the seller is told how many filters the buyer chose and whether or not the buyer has bought the product.

The payoff for each round is calculated as follows:

If the buyer buys the product:

$$\text{Payoff of the buyer} = \text{Value (of product drawn)} - \text{Price in €}$$

$$\text{Payoff of the seller} = \text{Price in €}$$

Therefore, if the buyer buys the product, he or she will receive as payoff the value of the product minus the price. In this case, the seller will receive a payoff that is the same amount of the price.

If the buyer does not buy the product:

Payoff of the buyer = 0

Payoff of the seller = 0

Therefore, if the buyer does not buy the product, both buyer and seller will each receive a payoff of 0 € in this round. The seller therefore only receives the price if the buyer actually buys the product.

You will find out at the end of the experiment which round will become payoff-relevant.

Do you have any questions, or is anything unclear? If yes, please raise your hand now. We will then be glad to assist you. You may continue reading the instructions on the next page.

Description of the Second Part

This part of the experiment consists again of 15 rounds. At the end of the experiment, exactly one round of this part will be chosen randomly by the computer for payoff. Since you will only be told at the end of the experiment which round is payoff-relevant, you should make your decision carefully in each round.

[*Only for first part buyers:* The structure of this part is identical to the first part and you keep your role as in the first part – you are again a buyer. However, in this part a computer algorithm takes over the role of the seller.

After setting a number of filters, a computer algorithm determines a product price. Be aware that the computer algorithm saves your choice of the filter number and can use this information for the calculation of the prices. The computer algorithm is programmed with the goal to maximize profits from selling. Additionally, the chosen price will be changed by a random value. The computer algorithm cannot use information from previous periods.]

[*Only for first part sellers:* The structure of this part is identical to the first part and you keep your role as in the first part – you are again a seller. However, in this part a computer algorithm takes over the role of the buyer.

Again, you determine a price for every possible filter choice of the buyer (0, 1, 2, 3, 4, or 5 filters). The computer algorithm receives a target product, as the buyer did before, which you cannot observe. Afterwards the buyer chooses the number of filters (0, 1, 2, 3, 4, 5) randomly, whereas every choice is equally likely. Then, the computer algorithm receives a randomly drawn product from the remaining products. He buys the product if the value for him is at least as high as the price that you set for the chosen filter amount.]

Do you have any questions, or is anything unclear? If yes, please raise your hand now. We will then be glad to assist you.