ROCKWOOL Foundation Berlin
Institute for the Economy and the Future of Work (RFBerlin)

DISCUSSION PAPER SERIES 59/25

Obfuscation in Competitive Markets

Ernst Fehr, Keyu Wu

www.rfberlin.com AUGUST 2025

Obfuscation in Competitive Markets

Authors

Ernst Fehr, Keyu Wu

Reference

JEL Codes:

Keywords:

Recommended Citation: Ernst Fehr, Keyu Wu (2025): Obfuscation in Competitive Markets. RFBerlin Discussion Paper No. 59/25

Access

Disclaimer

Imprint

RFBerlin
ROCKWOOL Foundation Berlin –
Institute for the Economy
and the Future of Work

Gormannstrasse 22, 10119 Berlin
Tel: +49 (0) 151 143 444 67
E-mail: info@rfberlin.com
Web: www.rfberlin.com

Obfuscation in Competitive Markets

Ernst Fehr **Keyu Wu**

Department of Economics

University of Zürich

30 June 2025

Revised Version

ABSTRACT

In many markets, firms increase product complexity through add-on features, which can make the evaluation and comparison of products difficult and thus increase buyers' search cost. Does this product obfuscation limit buyers' search behavior and induce them to buy overpriced products? And if so, why does competition not eliminate obfuscated products? We show – based on competitive experimental markets – that if add-ons merely complicate the products without generating additional surplus, obfuscation via product complexity indeed becomes fragile because buyers display an aversion against complex products. However, in the presence of surplus-enhancing add-ons, obfuscation via product complexity becomes a stable market feature that severely constrains the depth and breadth of buyers' search. Sellers anticipate and take advantage of this by making unfavorable product features less visible and selling add-ons persistently above marginal cost. Even the most favorably priced product in the market is offered above marginal cost, and buyers persistently fail to find the best product such that inferior products have a good chance of being bought, leading to enduring price dispersion. Surplus-enhancing obfuscation opportunities are the causal driver of persistent profits and price dispersion because if we remove these opportunities, overall prices quickly converge to marginal cost.

1 Introduction

Complex products and price schedules are a frequent feature of modern economies. Complexity, however, comes with its own problems because it makes products difficult to evaluate and compare between competing sellers. The literature (e. g., Ellison and Ellison 2018) has pointed out that any action by the sellers that raises potential customers' cost of price and attribute search can be thought of as obfuscation, and obfuscation through complicated product attributes appears widespread.[1] For example, the selling of cars, printers, and cameras is often combined with complex bundles of differently priced add-ons and accessories. Electronic products like personal computers and memory modules are often advertised at low prices for a very basic product with limited memory, low capacities, and short warranties, but offer many kinds of complex upgrades separately (Ellison and Ellison, 2009). Similar situations have been proliferating in other industries and marketplaces like markets for transport and tourist services, financial retail markets, insurance markets and educational markets (Muir, Seim and Vitorino 2013; Miravete 2013; Grubb 2015; Greenleaf et al. 2016; Célérier and Vallée 2017; Bhargava, Loewenstein and Snydor 2017; Seim, Vitorino and Muir 2017; Ellison and Ellison 2018; Richards et al. 2019). The potential harm of complex prices and junk fees has also attracted attention from regulators in many industries (Greenleaf et al. 2016; Bourne and Bagley, 2023).

The widespread existence of markets with complex product attributes raises a number of fundamental questions such as whether sellers can use complex product designs and pricing rules to increase their profits at the expense of consumers' welfare? Is, for example, the existence of markets with obfuscation due to a pre-existing lack of competition or can profit-increasing obfuscation persist even in extremely competitive markets that would – in the absence of obfuscation – approach marginal cost pricing? Can profit-increasing obfuscation persist in an environment in which consumers can acquire experience, or do consumers over time see through the veil of obfuscated products, and shy away from them, such that firms with transparent products have a competitive advantage? Can obfuscation opportunities alone cause a persistent violation of the law of one price? And finally, what are the welfare effects of markets with persistent obfuscation, i.e., do these markets reduce consumers' overall welfare relative to a perfect (Bertrand) competition yardstick? And if so, what are the underlying mechanisms?

An accumulating body of theoretical papers examine the mechanisms underlying persistent obfuscation and its consequences in markets (e.g., Ellison, 2005 and 2006; Spiegler, 2006; Gabaix and Laibson, 2006; Carlin, 2009; Wilson, 2010; Ellison and Wolitzky, 2012; Piccione and Spiegler, 2012; Chioveanu and Zhou, 2013; Clippel, Eliaz and Rozen 2014; Spiegler, 2016; Heidhues, Kőszegi and Murooka, 2017; Heidhues and Köszegi 2018; Hämälainen 2018; Shulman and Geng, 2019; Hefti, Liu and Schmutzler, 2020; Heidhues, Johnen, Köszegi 2021). The empirical literature has reported somewhat mixed results. Some papers have found evidence consistent with theories that predict that firms can increase profits through obfuscation while others find that shrouding products may have negative effects on firms.

For example, Ellison and Ellison (2009) document that charging a low price for a firm's low-quality product attracts many consumers and helps increase sales of medium and high-quality products. Chetty, Looney and Kroft (2009) show that demand elasticities are lower when tax increases are shrouded

[1] This characterization of obfuscation also means that obfuscation may occur even if sellers do not deliberately intend to confuse or mislead the buyers, i.e., sellers are not required to have an independent motive to mislead buyers. Rather, the increase in buyers' cost of price and attribute search may simply be a byproduct of attempts to maximize profits.

compared to when they are not shrouded. And Celerier and Vallee (2017) show that financial products with more attractive headline returns are associated with higher complexity and more risk for the customers but generate higher mark-ups for firms. Other literature has, however, suggested countervailing forces to obfuscation. For example, Gaudeul and Sugden (2012) proposed that consumers prefer transparency and simplicity while others have examined this question empirically (Xia and Monroe 2004; Crosetto and Gaudeul, 2012; Repetti, Roe and Gregory 2014; Seim, Vitorino and Muir 2017; Sugden and Zheng, 2018). Chiles (2017), for example, reported that shrouding surcharges decreased sellers' reputation and Hossain and Morgan (2006) and Brown, Hossain and Morgan (2010) find that shrouding surcharges does not improve or even decrease revenues.

In this paper we tackle the above-mentioned fundamental questions with the help of a laboratory experiment with real monetary incentives and a market environment with competing sellers and buyers. We believe that the enhanced controls offered by an incentivized laboratory experiment can provide empirical insights that are difficult to obtain with field data. Take, for example, the question whether profit-increasing obfuscation can persist even in competitive markets that would – in the absence of obfuscation – approach marginal cost pricing. With field data, it is typically not possible to observe counterfactual markets, i.e., what would happen in the absence of obfuscation, while in our lab experiments, we can exogenously vary obfuscation opportunities such that we can identify their *causal impact* on a wide range of market level phenomena and individual behaviors of both sellers and buyers.

Our experimental design approximates a competitive six-seller online shopping market in which obfuscation opportunities are implemented by giving sellers the option to add extra features to a basic product. Importantly, the add-on features can be offered by all sellers. Therefore, they do not constitute unique technological innovations that provide sellers with a competitive advantage, but they just represent the detailed values and surcharges in which the different products may differ. Sellers stipulate a transparent headline price for their base product while the separately listed values and prices of the products' add-on features are only revealed if buyers take the time to learn them. The buyers are completely free regarding how much time they want to invest in searching the market, but time spent searching is costly for them. This design captures typical properties of complex products where consumers must invest effort to find and understand the net benefits of the product's features. Sellers and buyers interact for 20 periods in the market so that they can acquire experience, learn, and respond to the behavior of other market participants.

We implemented treatments that varied the available obfuscation opportunities. In the "*surplus-enhancing treatment*", obfuscation takes the form of surplus-enhancing add-on features that increase the available surplus because consumers value the add-ons above their production cost. Add-on features that enhance surplus but also make the products complex are prevalent in many real-life situations, such as quality upgrades for electronic products, extra services like extended warranties or the right to cancel a reserved flight or hotel. In contrast, in the "*surplus-neutral treatment*" all the surplus is already in the base product and adding additional "features" does not change the available surplus but merely partitions the product's total price into more complicated segments that require buyers to spend extra time to understand. This type of obfuscation reflects situations where webpages and product descriptions are complicated and thus impose extra cost and effort on the buyers. This treatment captures a situation that is at the center of many theories of obfuscation where sellers' decision to add complexity is independent of the *overall* value and price of the offered product.

To identify the causal effects of obfuscation opportunities, we also implemented a "*no-obfuscation control treatment*". The key difference between treatments with and without obfuscation opportunities is that in the latter buyers have direct access to transparent information about the *overall* net value of all products. This treatment resembles an ideal form of policy intervention that summarizes the products in the simplest possible form (e.g., a quality-adjusted price) to enforce competition. The no-obfuscation control treatment allows us to identify the extent to which obfuscation opportunities – rather than other factors – are causing persistent positive profits, price dispersion, or lower the buyers' welfare.

To what extent do competitive markets with surplus-enhancing obfuscation opportunities enable sellers to increase their profits at the expense of buyers' welfare? We find that these markets are characterized by a high level of product complexity and that sellers appropriate a substantial share of the total surplus that buyers would have received under marginal cost pricing. Even in the long run, when average prices have settled at a stable level, sellers can appropriate roughly 1/3 of the total surplus because (i) even the best available product in the market (i.e., the one with the highest surplus for the buyers) is associated with a substantial mark-up of prices over marginal cost, and (ii) buyers' persistently fail to purchase the best available product in the market. Moreover, substantial dispersion in offered and traded (overall) prices prevails throughout the 20 market periods, violating the law of one price. Remarkably, this dispersion prevails despite the fact that (i) sellers do not have heterogeneous production technologies and (ii) buyers do not have heterogenous tastes. These findings contrast sharply with the control treatment where obfuscation opportunities are absent. Here, the market quickly converges to marginal-cost pricing after only a few periods and almost all buyers purchase the best available product in the market, which generates strong competition among the sellers that drives up the buyers' share of the surplus.

Why can sellers persistently earn a relatively high share of the trading surplus in markets with surplus-enhancing add-ons? The answer hinges on how buyers responded to the presence of high product complexity and how this shaped the sellers' incentives. We find that, on average, buyers' breadth of search (i.e., the number of products they "visit"[2]) as well as their depth of search (i.e., the share of within-product add-ons they examine) is rather limited. Specifically, in roughly 40% of the cases the buyers visit only one or fewer products and in roughly 40% of the cases buyers see less than 60% of the add-ons of the visited products. Thus, in many cases buyers exerted little or no competitive pressure on add-on prices.

However, we also find substantial heterogeneity in buyers' search because some of them tend to be "browsers" who visit many products in a superficial way while others are "studiers" who tend to visit only a few products but examine them in depth – a finding that is reminiscent of the model by Heidhues, Johnen, Köszegi (2021). Moreover, we collected an independent measure of each buyer's (i) time needed and (ii) their subjective mental cost[3] of assessing the overall net value of products with varying numbers of add-ons. We find that buyers for whom the time needed to compute the overall net value of products is more convex (as a function of the number of add-ons) focus their search relatively more on breadth rather than depth; likewise, buyers with higher mental costs are also more likely to be browsers rather than studiers. Thus, the heterogeneity in buyers' time and mental costs of assessing the overall net value helps us understand why we observe heterogeneous search behaviors.

[2] A product is considered as visited if the buyer examines at least one add-on.
[3] Mental costs exist because assessing the overall net value is effortful and thus, buyers may have a willingness to pay to avoid it.

Obviously, the buyers' limited depth of search means that they do not know the values and prices of many add-ons of visited product, and their limited breadth of search implies that they do not know the add-ons of non-visited products. As a consequence, one would expect that buyers respond on average rather sluggishly to a product's own add-on values and prices as well as to those of competing products – a conjecture that is strongly supported by our data.[4]

Do the sellers anticipate the rather limited breadth and depth of buyers' search behavior? We find indeed that the sellers are very much aware of this. Sellers believe, e.g., that buyers visit, on average, only 3.23 of the six available products in the market and examine only 3.26 features of products with 6 add-ons. This has important consequences for the sellers' incentives. Sellers have, in particular, an incentive to (i) place the worse product features on the list of add-ons last – because that maximizes the chances that buyers do not see them – and (ii) to charge non-competitive add-on prices. We find strong evidence that sellers indeed behave in this way. In fact, high non-competitive add-on prices are the key source of sellers' profits. Thus, the observed buyer and seller behaviors are consistent with rational responses to the prevailing incentives in the market.

Note, however, that sellers do not act as monopolists by extracting all the surplus generated from add-ons features, and add-on prices turn out to be persistently heterogeneous throughout the 20 market periods. This finding differs from several theoretical models that assume (e.g., Gabaix and Laibson, 2006; Heidhues, Kőszegi and Murooka, 2017; Heidhues, Johnen and Kőszegi, 2021) or predict (e.g., Diamond, 1979; Lal and Matutes, 1994; Ellison, 2005) no competition with regard to hidden product features which would enable firms to extract all the surplus from add-ons. As mentioned above, many buyers do visit two or more products in some depth which limits sellers' ability to set arbitrary add-on prices. In addition, the heterogeneity of buyers' search behavior allows for the persistent trading of products with heterogeneously priced add-ons.[5]

Moreover, to examine the robustness of our interpretation of buyers' and sellers' behaviors as rational responses to incentives, we also conducted further treatments in which we examined the behavioral responses to an *exogenous* increase in the pecuniary cost of search. If buyers and sellers indeed respond rationally to a search cost increase, we should observe that higher search costs

(i) reduce the depth as well as the breadth of search, and thus increase buyers' information imperfections. This, in turn, should further
(ii) reduce buyers' responsiveness to the changes in add-on prices and values of a given product as well as their responses to the add-on prices and values of the best competing product. As a consequence, these buyer responses
(iii) allow the sellers to enforce higher add-on prices such that they reap a higher share of the total surplus from the add-ons.

[4] We also find that buyers' purchase behavior is not affected by the number of add-ons per se or by the specific labels attached to the add-ons. Buyers appear exclusively interested in the pecuniary aspects (i.e., the overall value and price) of the products.

[5] The strong heterogeneity in add-on features may also have repercussions on competition that operates via headline prices. Dispersed "hidden" prices for add-ons means that headline prices provide only a very incomplete picture of the overall price of a good and there is always a need for buyers to search the add-on features to find better products. We find that more attractive headline prices are associated with less attractive add-on features, thus providing a potential reason for buyers to be suspicious of too low headline prices. However, the negative relation between headline prices and the attractiveness of add-on features is quite noisy which makes it difficult for the buyers to make reliable inferences about the quality of add-ons from headline prices.

We find indeed that all these predictions are strongly supported by the data.

Are sellers in markets with surplus-neutral obfuscation opportunities also able to persistently increase their profit at the expense of the buyers' welfare? Our data reveal a clear bi-modal pattern in this case: in half of the markets (that we label "high complexity markets"), product complexity is high and sellers are able to appropriate consumer welfare to a similar extent as in the surplus-enhancing treatment; but in the other half of the markets ("low complexity markets"), product complexity declines over time to rather low levels and buyers eventually receive almost all of the available trading surplus. An analysis of buyers' buying behavior in low complexity markets suggests that a potential reason for the decline under surplus-neutral obfuscation is that buyers shy away from buying complex products even after controlling for the products' values and prices. This aversion to complex products with surplus-neutral add-ons thus appears to be a force that mitigates individual sellers' incentives to offer complex products.

But why do buyers display complexity aversion in the markets with surplus-neutral add-ons but not in the market with surplus-enhancing add-ons? A possible reason appears to be that under surplus-neutral add-ons the overall net value of a product's add-on features is *negatively* related to the complexity of the product. Thus, more complex products are typically worse for the buyers. We also document that buyers notice this pattern, which provides a reason for shying away from complex products. In contrast, this relationship is *reversed* in the market with surplus-enhancing add-ons. Here, a product with a higher number of add-on features is typically a better product for the buyers. Therefore, there is little reason for complexity aversion to prevail in markets with surplus-enhancing add-ons, which contributes to the sustainability of the associated competition-mitigating consequences of complex products.

Our study contributes to the literature in behavioral industrial organization that studies the empirical functioning of markets in which consumers have limited attention or imperfect knowledge about product attributes. We believe that a main contribution of our paper consists in the identification of the causal impact of different types of obfuscation opportunities on a wide range of individual level behaviors and market level outcomes. Our paper is also related to the literature that suggests or hypothesizes that consumers value transparency and simplicity (Xia and Monroe 2004; Crosetto and Gaudeul, 2012; Gaudeul and Sugden, 2012; Repetti, Roe and Gregory 2014; Seim, Vitorino and Muir 2017; Sugden and Zheng, 2018). We document this desire for transparency (i.e., complexity aversion) by showing that in some of our markets the buyers shy away from more complex products that are otherwise identical, i.e., even after controlling for the products' overall values and prices. In addition, our findings inform us about whether and under which conditions complexity aversion can undermine the stability of obfuscation at the market level and drive the whole market towards transparency.

We believe that our empirical findings may also be of interest for the theoretical literature on obfuscation in markets (e.g., Ellison, 2005; Spiegler, 2006; Gabaix and Laibson, 2006; Carlin, 2009; Ellison and Wolitzky, 2012; Piccione and Spiegler, 2012; Chioveanu and Zhou, 2013; Heidhues, Kőszegi and Murooka, 2017; Shulman and Geng, 2019). In particular, in view of the paramount importance of the assumptions on buyer behavior in theories of obfuscation, we believe that our detailed empirical results on the heterogeneity of buyers' search patterns in terms of "browsing" and "studying", and the underlying time and mental costs of assessing the overall net value of products that drive these patterns, may be useful for future theorizing. However, unlike Heidhues, Johnen, Köszegi (2021) and many others (Ellison, 2005; Gabaix and Laibson, 2006; Heidhues, Kőszegi and Murooka, 2017), we find that competition via add-on prices, albeit strongly mitigated, is still present, and results in heterogeneously priced add-ons. Similarly, some theories of search and obfuscation predict that all buyers with positive search cost visit only one product (e.g., Stahl, 1989; Ellison and Wolitzky, 2012) which confers enormous market power to firms. This prediction contradicts our finding that a substantial share of the buyers visit two or more products which puts tighter limits on sellers' market power.

Our paper is also related to the literature in marketing that examines how partitioned prices or "drip pricing" affects consumers' perceptions, attitudes and behaviors (e.g., Morwitz, Greenleaf and Johnson 1998; Lee and Han 2002; Xia and Monroe 2004; Bertini and Wathieu 2008; Völckner 2012; Robbert and Roth 2014; Robbert 2015; Greenleaf et al., 2016; Dallas and Morwitz, 2020). Our study differs from this literature by explicitly embedding buyer-seller interactions into a competitive market environment. This makes it possible to study the implications of buyer behaviors for market level phenomena – such as positive profit shares or persistent price dispersion – and to study the mechanisms through which these phenomena are generated because we can explicitly observe the actions of buyers *and* sellers.

Finally, our paper also contributes to the experimental literature related to obfuscation and bounded rationality in consumer behavior. One part of this literature studies *boundedly rational individual* decision-making in complex situations such as in a complex product space (Crosetto & Gaudeul 2011; Jin, Luca and Martin 2015; Kalayci & Serra-Garcia 2016; Sugden and Zheng 2018; de Clippel and Rozen, 2021). These experiments neither allow for interactions in markets nor is there competition among sellers. The other part of the literature studies obfuscation in market environments but in many of these papers one side of the market in the treatment or the control condition is represented by virtual or robot agents (Kalayci and Potters, 2011; Kalayci, 2015, Gu and Wenzel 2015; Kalayci 2016; Crosetto and Gaudeul 2017). Our experiment differs in many dimensions from these papers but perhaps the biggest difference is that none of these papers studies the buyers' endogenous search behaviors. This means that no information exists, for example, on the depth and breadth of buyers' search, implying that the degree of buyers' incomplete information, and how sellers respond to it, remains unknown.

The rest of the paper is organized as follows. In Section 2, we describe the design of the market with obfuscation opportunities and the treatment variations that we implemented. Section 3 reports our detailed empirical findings in the market with surplus-enhancing obfuscation opportunities. After this we contrast our findings in the surplus-enhancing treatment with those of the surplus-neutral treatment in section 4. Section 5 concludes the paper and provides some additional discussions.

2 Experimental Design

To identify the causal impact of obfuscation on market outcomes, we compare a competitive market with obfuscation opportunities ("OO Market") with a control treatment that is identical in every respect except that the sellers can no longer obfuscate their products (market with no obfuscation, denoted as "NO Market"), i.e., the buyers can fully understand the characteristics of the offered products with ease. In our experimental design – described in more detail below – sellers are given the opportunity to add extra features (i.e., add-ons) to their basic product; each extra feature has a well-defined objective value for the buyers, and the price of each extra feature is also available for the buyers to observe. But because there are potentially many extra features, the buyers may have to invest some time to view them, compute the overall value of a product, and compare the products in the market, which captures a general characteristic of markets with complex products. In contrast, obfuscation is removed in the control treatment by informing the buyers transparently about the *overall* objective (pecuniary) value of *each* product in the market, which renders the market very transparent.

An experimental session consisted of three parts. In Part 1, participants interacted in a market with obfuscation opportunities (OO Market). In Part 2, the same participants interacted in a market without obfuscation opportunities (NO Market).[6] In Part 3, we collected additional measures of the extent to which complex products with a varying number of add-ons affect the buyers' time and mental cost of assessing the overall net value of products; these measures may help us better understand the mechanisms through which obfuscation works. During the experiment buyers and sellers can earn money in the form of experimental currency units (ECUs) that are exchanged into real money at the end of the experiment according to a publicly known exchange rate.

2.1 The market with obfuscation opportunities (OO Market)

In a market, 16 participants interact in a posted-offer institution for 20 trading periods. Among the 16 participants, 6 of them are randomly assigned to be sellers and the other 10 are assigned to be buyers throughout the entire 20 periods. The sellers and buyers trade experimental goods that are labelled as phones. Each phone consists of a basic phone and some extra features. Sellers can offer basic phones by incurring the same marginal production cost of 5 Experimental Currency Units (ECU), while buyers' valuations of a basic phone differ. In each period, a buyer's valuation of a *basic* phone is a random value out of five possible values: 0, 5, 10, 15, or 20 ECU; each of the five possible values is randomly assigned to two of the buyers in any given trading period. In each trading period, sellers can set a base price for their basic phone. Each buyer can buy a maximum of 1 unit, but each individual seller can serve the demands of all buyers. As 6 sellers are competing to serve the demands of 10 buyers with no limits in supply, we ensure that the market is very competitive.

On top of the basic phone, sellers can also add extra features to their products. Figure 1 shows an example of how these products are presented in the experiment. Each extra feature has a label, a value, and a price. Each feature provides buyers of this product with the corresponding additional value on top of their own basic values, but generates an additional production cost for the seller. The add-on prices, on the other hand, represent additional charges for the respective add-on feature but they can also be interpreted as representing other negative product attributes. This design maps real-life situations in which firms add and separately list many features to their product, like fancy technical properties, quality upgrades, extra services, extended warranties, express shipping, etc. While these features are often indeed valuable for consumers, they also increase the product's complexity and make the comparison of products difficult because their characteristics and prices differ in many dimensions. For simplicity, buyers of any product need to buy all the extra features that come with the product. This design captures many situations in naturally occurring markets.[7] In addition, to keep the set-up simple, a seller can offer only one contract for his product, which rules out the screening of buyer types.

[6] To examine whether there are spillover effects from the OO Market (Part 1) to the NO Market (Part 2) we also conducted control sessions in which the NO Market took place without a preceding OO Market. It turns out that the market outcomes in the NO Market are similar regardless of whether there was a preceding OO Market or not.
[7] This design feature differs from the assumptions of some theoretical papers (e.g., Gabaix and Laibson 2006). However, it generates strong incentives to examine all add-ons which may reduce sellers' ability to sell bad products to the buyers. In addition, this feature captures aspects of naturally occurring markets. To see this, consider, first, that in many cases the listed features and add-ons are indeed inseparable components of the product,

In this setting, we can denote buyers' basic value as v_b, the base price as p_b, and sellers' marginal cost of producing a basic phone as c_b. For a product with n extra features, suppose the value and price of the i^{th} feature is v_f^i and p_f^i respectively, then we can denote $v_f = \sum_{i=1}^{n} v_f^i$ as the aggregate feature value for a product with n extra features, $c_f(v_f)$ as the aggregate cost of producing v_f and $p_f = \sum_{i=1}^{n} p_f^i$ as the aggregate feature price. Then, from the sellers' perspective, their profits per unit sold, π^S, is given by their earnings from both the basic phone π_b^S and the extra features π_f^S, which can be calculated as:

$$\pi^S = \pi_b^S + \pi_f^S = (p_b - c_b) + (p_f - c_f(v_f))$$

From the buyers' perspective, their total earnings from buying a product, π^B, is given by the earnings from both the basic phone π_b^B and the extra features π_f^B, which can be calculated as:

$$\pi^B = \pi_b^B + \pi_f^B = (v_b - p_b) + (v_f - p_f).$$

Thus, buyers who maximize their total earnings π^B, have homogenous preferences with regard to all products in the market. In addition, $\pi^S + \pi^B$ provides a measure of the total surplus from a trade.

The Seller Stage

In the OO markets there are 20 trading periods, and each trading period consists of 3 stages: "Seller Stage", "Buyer Stage", and "Feedback Stage". In the "Seller Stage", sellers decide which product they want to offer. To simplify the sellers' decisions, we do not require them to fix a separate value and a separate price for each extra feature they provide. Rather, in addition to setting a non-negative base price, they only need to determine the number of features n, the aggregate feature value, v_f, and the aggregate feature price, p_f. Given the chosen levels of v_f, p_f, and n, the computer then randomly assigns a one-digit number to each feature value and feature price so that the two sums, v_f and p_f, are exactly met. Sellers can also re-randomize the computer's assignment of numbers as many times as they want until they are satisfied with how their products look.

We intentionally restricted the feature values and feature prices to single-digit numbers. This way, products are not too complex.[8] To further keep product complexity within limits, we also restricted the

such as many "features" in bundle sales and composite products like camera packages or insurance products. Likewise, printers or phones may be designed in such a way that consumers can later only use the ink or the earphones that are compatible with the original products. Second, buyers often already have a product with certain features in mind that they would like to buy before a purchase. For example, when consumers book flight tickets, they may already know that they need the right to change their bookings in the future or how many bags they want to check in. Therefore, even though these features appear optional, consumers in fact just need to find the full price that indeed contains all these features. Third, even when add-ons are actually optional, firms often design the descriptions of the products so that the basic product looks significantly worse than the upgraded product with add-ons (Ellison and Ellison, 2019), which entices many consumers to actually buy the add-ons. Finally, note that different firms may indeed offer different add-ons in our experiment. Thus, by deciding from which firm to buy, the buyer has still some discretion about the add-on he or she is buying.

[8] It would have been easy to introduce additional forms of complexity, for example, by allowing for buyers' uncertainty about the feature's values, or making some add-ons complements or substitutes for each other. However, if we find that already our "mild" obfuscation opportunities enable the sellers to appropriate substantial rents, then the more complex forms of naturally occurring obfuscation opportunities can be expected to cause much worse outcomes for the consumers. Note also that the single-digit constraint can also be interpreted as a

maximum number of features to 6. Furthermore, a seller of a product with $n \geq 1$ features could maximally add an aggregate feature value v_f of up to $7n + 2$. This upper bound on the aggregate feature value ensures that the single-digit constraint for individual extra features can be met. To keep things simple, the sum of extra prices is subject to the same aggregate constraint.[9] Note, that this design gives sellers the opportunity to make their products more complicated than the number of add-ons that are needed to generate the desired aggregate feature value v_f. For example, a seller who wants to provide add-on features with an aggregate value of, say v_f = 29, can do so with 4, 5, or 6 extra features.

Note that the extra features of the products in this experimental design are not any seller's unique technological innovations that provide *per se* a competitive advantage relative to the other seller's products. Rather, the possibility of adding extra features is a tool that can be used to generate additional value for the buyers, and all sellers have the same opportunities to use this tool. In addition, since all buyers derive exactly the same value from a given extra feature with no uncertainty or noise, it is very likely that buyers do not have a pre-determined taste or preference over certain types of products; rather, they just want to find the product that provides the highest monetary payoff π^B, a conjecture that we will later verify on the basis of buyers' observed behavior. Therefore, if competition is fully at work and buyers choose only those products that give them the highest overall earnings π^B, the extra features should not enable sellers to earn positive profits.

After all the 6 sellers have determined their products for the current trading period, the "Buyer Stage" starts, in which all the 6 products are displayed to buyers, in an order that is completely randomized across periods. Because sellers and buyers do not have an ID and remain anonymous to each other, our design eliminates any reputation concern. Sellers' reputation may of course play a role in naturally occurring markets with obfuscation opportunities, but reputation formation may interfere with obfuscation in multiple conflicting ways[10]. These confounds may make it hard to draw clean inferences about obfuscation behavior in the field, and most theoretical models of obfuscation also do not yet consider reputation formation. Therefore, here we are – as a first step – interested in how such a market operates in the absence of reputation formation opportunities.

The Buyer Stage

At the beginning of the "Buyer Stage", buyers are informed of their basic values in the current period, and then they can start shopping. When buyers are shopping, their time is valuable – for every second buyers spend in the market before they make a decision, they incur a cost of 0.1 ECU. The rationale for this is that, in reality, consumers typically have considerable opportunity cost of time from searching the market. In the experiment such time costs of search are absent because the subjects have already committed to participate in the experiment, i.e., they have no alternative use for their time during the

technological constraint that puts an upper bound on the value that each extra feature can provide. Such a constraint is plausible because it is generally not possible to generate arbitrarily high values without adding more features to a product.

[9] To further ensure that sellers understand the decisions they make, before a seller commits to an offer, he or she can try out many decisions; for each decision, the computer automatically calculates the overall cost of extra features c_f, earnings from extra features π_f^S, and profits per unit π^S. Sellers are also required to answer a list of comprehension questions correctly and go through a practice round before the actual market interactions start.

[10] For example, if buyers dislike obfuscation, firms may avoid obfuscation for reputational reasons. Or if buyers want to save search costs by displaying brand loyalty and buying from the same firm over time, firms may have additional sources of positive profits because of buyers' limited willingness to switch.

experiments. For this reason, we implemented an explicit, yet small, monetary cost of searching the market. On the buyers' screen (see Figure A1 in Appendix 1 for example screens), they first see only the base prices of the 6 products while the extra features of the various products are not immediately visible. However, the buyers can inform themselves about a product's extra features by "visiting" it.

To "visit" a product, a buyer has to click on the product on the screen. One click makes its first feature appear, the next click makes the next feature appear, and so on. Moreover, any product will only be clickable every 2 seconds. This design mimics typical situations of product search in that consumers always need some time to find the next piece of useful information about a product; moreover, upon seeing a product feature, consumers always need time to understand it and figure out its subjective value. Thus, compared to such time-consuming search in real online markets, the 2-second delay is rather conservative.[11] In fact, the time cost of 0.1 ECU per second is also chosen in such a way that the buyers can acquire full information about all products and still earn a substantial profit[12].

Overall, our design allows us to obtain a detailed dataset on which products buyers examine, how many add-ons they examine, and the order and overall duration of search. Note also that there is no limit on the total time that any buyer can spend in the market, i.e., they were free to search the market as long as they wanted. Thus, if buyers are imperfectly informed about the available product properties and add-on prices, they voluntarily forgo additional information. Buyers were also not forced to buy any of the available products, i.e., they could completely refrain from trading if they wanted.

Information conditions

All market participants know that there are 6 sellers and 10 buyers. They also know that the products consist of a basic phone plus a variable number of add-on features that are determined by the sellers. For both buyers and sellers, the potential characteristics of products are explained with the help Figure 1 above.

The sellers also know that the buyers' valuations for the basic phone differ between buyers and across periods and are in the range between 0 and 20 ECUs. We provided this upper bound of the valuation to avoid that sellers charge unreasonably high prices that could not be afforded by any buyer. The sellers also know that the buyers in the OO Markets see initially only the base prices of the six basic phones but that they can acquire information about the add-ons of each product sequentially by clicking repeatedly on the respective product. This was communicated by showing the sellers screens like those in Figures A1a – A1c in Appendix 1 in the experimental instructions.

After all buyers made their buying decisions, the trading period proceeds to the "Feedback Stage" (see Figure A2 in Appendix 1 for an example screen). In this stage, the sellers receive feedback on the details of the 6 products that were offered in the current period: their respective base prices, and the values and

---

[11] Moreover, only the features of the currently visited product are visible on the buyer's screen; that is, when buyers switch the product they visit, the features of the previous product disappear. This design approximates situations in which consumers search sequentially, i.e., where they do not have simultaneous access to all the available information. Our buyers can, however, use paper, pen and blank spaces on the shopping screen to record their findings from searching the add-ons.

[12] To mitigate "psychological liquidity constraints", we endowed each buyer with 8 ECU in every period which provides a "time budget" of 80 seconds. This is far above the actual average search time of 25.7 seconds that subjects actually spent in the market, i.e., buyers have enough liquidity to finance their search cost. Even if all 6 products had the maximal number of add-ons and a buyer would want to see them all, the endowment of 8 ECU would suffice to render all add-ons visible.

prices of each extra feature of all products. In addition, they are informed about relevant summary statistics associated with each product: the aggregate feature value v_f, the aggregate feature price p_f, the aggregate feature cost c_f, earnings from extra features ($p_f - c_f$), and profits per unit sold π^S. A seller also privately sees how many units he sold and his realized total earnings in the current period. The feedback for sellers provides them with substantial information about the whole market and thus may enhance competition. Buyers, on the other hand, are shown their realized earnings from trading, total time cost, and total earnings in the current period.

2.2 The market without obfuscation opportunities (NO Market)

In this treatment, the participants have the same role assignments and interact with the same group of people under the same conditions as the OO Market except that the computer provides information during the "Buyer Stage" that makes each product's overall net value immediately visible for each buyer.

Recall that the buyers' earnings from a trade π^B are given by,

$$\pi^B = \pi_b^B + \pi_f^B = (v_b - p_b) + (v_f - p_f) = v_b + (v_f - p_b - p_f).$$

Note that the buyers' basic values, v_b, which are randomly assigned and privately communicated to the buyers at the beginning of every period, do not affect the relative attractiveness of the different products because they refer to the value of a basic phone (which is, for any given buyer, identical across products). In contrast, the base price, p_b, the aggregate feature value, v_f, and the aggregate feature price, p_f, all depend on sellers' decisions. They typically vary across products and can be summarized by

$$v_o \equiv v_f - p_b - p_f.$$

Therefore, buyers who maximize π^B only need to know the "overall net value" v_o to assess the relative attractiveness of the available products, and the computer publicly provides this information to the buyers in the NO Market. In other words, by providing information about v_o for all available products, π^B-maximizing buyers no longer need to examine the products' add-ons, so that the need for search is essentially reduced to zero. The sellers in the NO Markets also know that the buyers are informed about v_o for each of the six available products. Therefore, we conjectured that this market will relatively quickly converge to the competitive equilibrium in which the maximum surplus is produced and appropriated by the buyers. For this reason, the NO Market lasted only 10 periods.

The NO Market mimics an ideal form of policy intervention that requires all the sellers to summarize the useful information of their products in one quality-adjusted price (or in a way as simple as possible). For example, personal loan providers are required to summarize their products by just one Annual Percentage Rate (Ellison and Ellison, 2020). This treatment is virtually similar to the policy intervention proposed in Ellison (2005): sellers are required to advertise one price and provide all the add-ons free of charge[13]. This way, products' overall net values are transparent to buyers.

[13] Ellison (2005) also shows theoretically that even when the add-ons are optional and consumers have heterogeneous preferences over add-ons, this policy intervention benefits both consumers who buy the add-on *and* consumers who do not buy the add-on.

2.3 Further treatments conditions and robustness checks

We implemented three treatments that enable us to examine how markets vary in response to different forms of obfuscation opportunities (see first 3 rows in Table 1 below). In addition, we implemented two further treatments that allow us to study how buyers' search and buying behavior is causally affected by an exogenous increase in search costs and whether sellers can take advantage of this by reaping a higher share of the surplus from add-ons.

We conducted three main treatments and two "robustness check treatments" (see Table 1). The main treatments differ according to the cost function $c_f(v_f)$. In the Half-Cost Treatment (HCT), the cost of producing extra features, c_f, is always 50% of the aggregate features values , v_f, so that adding more feature values increases the available total surplus but typically also a product's complexity (i.e., the number of add-ons). While this type of obfuscation by producing surplus-enhancing add-ons likely approximates a frequent form of obfuscation in naturally occurring environments (e.g., upgrades of products, extra accessories, faster shipping, etc.), we also want to understand how much obfuscation occurs in a competitive market when the marginal surplus from extra feature values can become negative. In this setting, sellers may even produce "socially wasteful" add-ons for which they can charge nevertheless positive prices (Heidhues and Kőszegi, 2017). Therefore, we implemented a second between-subjects treatment where the cost of producing extra features, c_f, is a convex function of the aggregate feature values, v_f, so that adding additional feature values beyond v_f = 30 reduces the total surplus. We call this treatment "Convex Cost Treatment (CCT)". This treatment enables us to examine whether sellers even offer inefficiently high levels of v_f (which are typically associated with higher complexity).

Since we are also interested in the effects of obfuscation technologies that merely increase complexity without generating any additional surplus, we implemented a "Surplus-Neutral Treatment (SNT)" where the cost of producing add-ons is always identical to the add-ons' values, i. e., $c_f = v_f$. Unlike the two previous treatments in which sellers may implement add-ons to provide additional surplus, in the SNT all the surplus is already in the basic product, i.e., for any level of aggregate feature values v_f the surplus remains the same.[14] Therefore, a seller who does not want to obfuscate can merely charge a base price without any add-ons (i. e., v_f = 0) and no surplus will be sacrificed. Any addition of extra features only has the effect of making the product more complex, which requires buyers to spend more time examining the product and increases the buyers' search costs. This case where sellers' decision to complicate the product is independent of the product's pecuniary aspects has been frequently examined in theoretical work. In naturally occurring markets, this corresponds to situations where sellers may, e.g., design their websites so that it is not straightforward to discover additional charges, or they may use lengthy descriptions and difficult language that make it difficult to discover a product's unfavorable attributes.

From the sellers' viewpoint, adding extra features may make sense if sellers believe that a more complex product enables them to reap higher profits. This could be the case if they believe that buyers have a preference for more add-ons per se, or that buyers may engage only in limited product search and do not thoroughly investigate all add-ons. By analyzing buyers' behavior and sellers' beliefs about buyers' behavior, we can distinguish between these two potential reasons for providing add-ons.

[14] Experimentally, the total available surplus in this treatment is kept approximately identical to the actual surplus in the HCT and CCT by moving all the surplus to buyers' basic values.

Table 1: Treatment Conditions

	Main Treatments	
	Market with obfuscation (OO Market)	**Market with no obfuscation (NO Market)** [15]
Half Cost Treatment (HCT)	6 markets; 20 periods	6 markets; 10 periods
Convex Cost Treatment (CCT)	5 markets; 20 periods	5 markets; 10 periods
Surplus-Neutral Treatment (SNT)	6 markets; 20 periods	6 markets; 10 periods
	Robustness Check Treatments	
Convex Cost Treatment (CCT) with low search cost	4 markets; 30 periods; low search costs	4 markets; 10 periods; with all add-on features shown
Convex Cost Treatment (CCT) with high search cost	4 markets; 30 periods; high search costs	4 markets; 10 periods; with all add-on features shown

Note: In the HCT, add-ons are always surplus-enhancing; in the CCT, they are surplus-enhancing within a certain value range and in the SNT, add-ons never increase a product's overall surplus. In the robustness check treatments, the OO Market took place for 30 periods, which enables us to study whether individuals' behavior and market outcomes still adjust significantly after 20 periods. Appendix 6 shows that this is not the case. The NO Market was implemented in two versions. At the buyer stage of the main treatments, the computer displayed only the overall net value v_o of all six products (see buyer screen in Figure A3). Under the assumption that buyers only care for v_o, which we verify in the results sections below, v_o alone gives buyers all they need to make a rational decision. In the NO Markets of the robustness check sessions, the buyers saw – in addition to v_o – the add-ons of all products (see buyer screen in Figure A4 of Appendix 1). It turns out (see Appendix 7) that buyer behavior and convergence to competitive outcomes is identical across these two versions of the NO Markets.

Note also that buyers in all of our treatments only know that sellers incur a cost when they add extra features, but they are never informed about the sellers' precise cost levels. We implemented this for external validity reasons because buyers in naturally occurring markets almost never know the sellers' cost situation.

Finally, we implemented two additional treatments (labelled "robustness check treatments") that served the following purpose. First, we varied the buyers' monetary cost of search by implementing a low and a high search cost treatment. This allows us to examine how the market participants and market outcomes respond to an increase in search cost. Second, we ran the markets for 30 periods, which enables us to check whether market convergence requires more than 20 periods. And third, we implemented a different variant of the NO Market where – in addition to the overall net values (v_0) of the offered products – all add-on features of the offered products are also immediately visible to buyers so that only the need for search is removed relative to the OO market. Thus, by comparing this variant of the NO Markets with the previous variant, we can examine whether behavior and, in particular, convergence to competitive outcomes, is affected by providing detailed information about the products' add-ons in addition to products' v_0. We document the robustness of our results with regard to (i) an increase in the number of periods and (ii) detailed information about add-ons in the NO market in Appendix 6 and 7.

[15] In each of the treatment conditions listed in Table 1, the NO Markets took place after the OO Markets. However, to control for potential spillover effects of the OO Market on the NO Market, we also conducted a few NO Markets (not shown in Table 1) without a preceding OO Market. The market outcomes in the NO Market turned out to be very similar regardless of whether there was a preceding OO Market or not.

2.4 Additional measures, procedures, and subjects

At the end of the OO Market in Part 1, after the final trading period, we also elicited buyers' beliefs about the buyers' average aggregate net feature value ($v_f - p_f$) of products in the market in which they just participated. We elicited $v_f - p_f$ for each possible number of add-ons.[16] These estimates will allow us to understand whether the buyers perceived more complex products on average as more or less valuable. In several sessions buyers and sellers also answered a questionnaire that elicits (i) the sellers' beliefs about the buyers' search and demand behaviors, and (ii) the buyers' motives underlying their search and demand behavior, respectively.

In Part 3 of an experimental session, we collected additional measures that help us better understand how product complexity influences (i) buyers' time needed for assessing the overall value of a product and (ii) their purely mental costs of assessing the overall net value of products. Assessing the value of a product with more extra features is naturally associated with a higher cognitive effort. Therefore, the *per-add-on* time needed for assessing the overall net value may increase with the number of add-ons such that the total time neededfor assessing the overall net value may be a convex function of the number of add-ons. We implemented a task in Part 3 that enables us to study this question. In addition, we implemented a further task that allows us to measure the purely mental costs of assessing the aggregate value of add-ons. The details of these tasks are described in Appendix 2. At the end of the experiment, 10 random trading periods of the OO Market, 5 random trading periods of the NO Market, and earnings from Part 3 are paid out. This payment scheme alleviates House Money Effects – because subjects do not know how much money they accrue over time – and generates, at the same time, reasonable incentives to motivate subjects to make careful decisions in every period of the market.

The experiment was computerized and programmed with the experimental software z-Tree (Fischbacher, 1999). There are a total of 272 subjects in the first three session types of Table 1 (HCT, CCT and SNT) and 128 subjects in the CCT sessions with low and high search costs. Each experimental session lasted approximately 2.5 hours and was held in the Econ Lab at the University of Zurich. Each subject earned on average 65 Swiss Francs (CHF 65 ~ USD \$72). All subjects were recruited from the joint subject pool of the University of Zurich (UZH) and the Swiss Federal Institute of Technology Zurich (ETH). All interactions in the experiment were anonymous.

2.5 Discussion of potential outcomes

No Obfuscation (NO) Market

In a market with Bertrand competition prices should converge to marginal cost because as long as the lowest price is above marginal cost, a seller has an incentive to undercut that price and serve the demands of the entire market, while all other sellers earn only zero profits. Experimental studies on Bertrand markets show that the Bertrand outcome is typically quickly reached when there are 4 or more sellers (Dufwenberg and Gneezy, 2000; Huck, Normann and Oechssler, 2004). If buyers in our NO Market care mainly about the overall net values v_o of the products the NO Market in our experiment is very similar to Bertrand competition. As we have 6 sellers, and each of whom can serve the whole market, we expect fierce competition in the NO Market such that sellers produce the maximum possible surplus,

[16] They were rewarded 5 ECU for each estimate within ±2 ECU of the actual average aggregate net feature value.

which is appropriated by the buyers. This also means that the overall price that trading buyers pay for a product (i.e., the base price plus the aggregate feature price) is identical, i.e., the law of one price holds.

Market with obfuscation opportunities (OO Market)

In contrast to the NO Market, buyers may not be able to assess and compare values and prices easily in an OO Market with complex products because they need to invest time and effort to understand each of the offered products.

Diamond (1971) first formalized a model of markets with exogenous search costs and theoretically showed that monopoly pricing that extracts all the surplus from buyers is an equilibrium in this environment.[17] Monopoly pricing outcomes are also predicted for products' hidden features in many theories of obfuscation and shrouded attributes under various other assumptions that imply limited competition regarding "hidden" add-ons and surcharges (Lal and Matutes, 1994; Ellison, 2005; Gabaix and Laibson, 2006; Heidhues, Kőszegi and Murooka, 2017; Heidhues, Johnen and Kőszegi, 2021).

However, the monopoly pricing result of Diamond (1971) is also fragile because it can unravel with an arbitrarily small inducement to visit multiple firms (Heidhues, Johnen and Kőszegi, 2021). Such situations are likely to occur when consumers have a heterogeneous willingness to search across products' add-on features, and firms may face some incentive to undercut high add-on prices to attract the searching consumers. Thus, heterogeneous buyer search behavior may act as a constraint on the "overpricing" of add-on features, and sellers' who target different buyer types could also generate differentiation and dispersion in the prices of add-on features and lead to the violation of the "law of one price" (Carlin, 2009; Ellison and Wolitzky, 2012; Chioveanu and Zhou, 2013).

In the end, the final price level and price dispersion are likely to depend on how much competitive pressure consumers exert on sellers with their endogenous search activities and purchase behaviors. In our experiment, we are able to explicitly examine the breadth and depth of buyers' search, thus enabling us to assess how their search patterns differ from those postulated in different theories and what market consequences this may have. Moreover, we can measure the ways in which sellers respond to buyers' search and purchase patterns, thus providing empirical insights into the mechanisms through which obfuscation opportunities mitigate competition and cause a redistribution of the surplus from trade.

Both in reality (e.g., in online shopping markets) and in our experiment, the products' base prices are typically considerably more salient and transparent compared to the prices and values of the add-on features. This then means that base prices may be subject to stronger competition than add-on prices – an assumption or a prediction that almost all theories of add-on pricing make. However, stronger competition for base/headline prices does not necessarily mean that this competition is unconstrained. In a market with obfuscation opportunities, when buyers have a limited willingness to search, they may try to infer the hidden net values of add-on features from headline prices, especially if they experience heterogenous add-on features during their search. And if they sometimes have bad experiences with a product that had a low headline price[18], they may become suspicious of overly low headline prices, a possibility that is, e.g., discussed in Völckner, Rühle and Spann (2012), Heidhues, Kőszegi and Murooka

[17] Intuitively, this result follows because in equilibrium, all consumers rationally expect all firms to charge monopoly prices and have, therefore, no incentive to search at all and just buy one of the products randomly. This consumer behavior, in turn, gives firms no incentive to lower their prices below the monopoly level.

[18] Celerier and Vallee (2017) show, e.g., that financial products with more attractive headline returns are associated with higher complexity and more risks, an empirical regularity that may provide a rational basis for being suspicious of attractive headline returns.

(2017) and Shulman and Geng (2019). If buyers are indeed suspicious of low headline prices, then sellers may hesitate to lower their base prices. Because we can directly examine both the characteristics of sellers' offered products and buyers' search behavior in response to different headline prices, our data can shed light on these issues.

Finally, there is one aspect that has not yet received much attention in the theoretical literature on obfuscation – the possibility that consumers may be averse to product complexity. Some literature (Xia and Monroe 2004; Gaudeul and Sugden, 2012; Crosetto and Gaudeul, 2012; Repetti, Roe and Gregory 2014; Chiles, 2017; Seim, Vitorino and Muir, 2017; Seim, Vitorino and Muir 2017; Sugden and Zheng, 2018) has suggested that consumers appear to be averse to hidden fees and value transparency. Complexity aversion may be the direct consequence of the search cost that complex, obfuscated products impose on consumers (i.e., consumers may resent searching through obfuscated products) or it may result from consumers' experience that more complex products are, on average, more likely be associated with a "worse deal". Our experiment offers the opportunity to examine whether complexity aversion (i.e., whether buyers shy away from buying more complex products that are otherwise identical to less complex products) emerges in highly competitive markets and whether this aversion is sufficiently strong to affect aggregate obfuscation levels and potentially other market level outcomes.

3 The impact of surplus-enhancing obfuscation opportunities

In this section, we present the results for the two main treatments with surplus-enhancing add-on features (Half Cost Treatment and Convex Cost Treatment).[19] In Section 3.1 we examine aggregate market outcomes in terms of (i) buyers' surplus per trade, (ii) the dispersion in buyers' surplus, and (iii) the overall complexity of the products in the market. Then, we examine the micro-foundations for these market outcomes in terms of buyers' search and purchase behavior in Section 3.2. The findings in this section then enable us to examine and understand sellers' behavior (Section 3.3) and how the interactions between buyers and sellers generated the overall market outcomes in the OO Market. Finally, we will corroborate our findings with the results from robustness checks in Section 3.4.

3.1 Market outcomes

Our first primary result concerns how much of the total surplus buyers obtain on average in the market with no obfuscation (NO Market) and in the market with obfuscation opportunities (OO Market), and the extent to which the traded surplus is dispersed across buyers. We summarize the corresponding findings in:

Result 1 (Buyers' Surplus in OO and NO Markets):

(a) In the absence of obfuscation opportunities, buyers receive almost all the surplus available in the market. After the first few periods, buyers' surplus quickly converges to 97% of the total surplus.

(b) In contrast, buyers get a much smaller share of total surplus, initially as low as 11%, in the presence of obfuscation opportunities. Buyers' surplus converges towards roughly two thirds of the total surplus in the long run.

[19] We pool the data together for results where the two treatments differ only in irrelevant ways.

(c) Dispersion in individual buyers' surplus quickly becomes negligible in the NO Market, whereas large and stable dispersion always prevails in the OO Market, indicating a violation of the law of one price.

We document Result 1 in terms of the share of the traded surplus that buyers obtain on average in the market in percent of the maximally possible total surplus. Recall that buyers' earnings from a product are $\pi^B = (v_b - p_b) + (v_f - p_f)$. If we denote the aggregate feature value v_f that maximizes the total surplus from feature values by $v_f^{max} \equiv argmax\,[v_f - c_f(v_f)]$, the percentage of the total available surplus a buyer receives from a trade is given by

$$\frac{\pi^B}{\pi^B + \pi^S} = \frac{(v_b - p_b) + (v_f - p_f)}{(v_b - c_b) + (v_f^{max} - c_f(v_f^{max}))}$$

If at least one seller in the market provides the efficient level of extra features v_f^{max} and prices his/her product at marginal cost (i.e., $p_b + p_f = c_b + c_f(v_f^{max})$), then any buyer who buys this product earns the maximum possible surplus.

Figure 2 shows the average buyer surplus per trade in percent of the maximum possible total surplus over the course of the experiment in the NO and the OO Markets. In addition, the figure displays the average within-period market-level dispersion in buyer surplus with "deviation bars" that indicate plus/minus one standard deviation of the buyer surplus from the mean. In the NO Market, where search is removed by design, the average traded buyer surplus starts off at a very high level (83%), quickly increases to 94% in period 3, and finally reaches on average 97% of the total surplus in periods 6-10.[20] In addition, the dispersion of the buyer surplus very quickly becomes extremely small. Therefore, competition pushes the total surplus to its most efficient level and prices very close to marginal cost in the NO Market such that the law of one (overall) price holds.

In contrast, Figure 2 shows that buyers' surplus in the market with obfuscation opportunities (OO Market) is only 11% of the maximal total surplus at the beginning. Competition pushes the buyer surplus slowly up over time, and the buyers' share stabilizes at roughly 2/3 of the total surplus from period 13 onwards. The buyers' share of the surplus in the OO Market is thus significantly lower than the buyers' share in the NO Market (p = 0.000, t-test with standard errors clustered at the market level)[21]. These facts sharply contrast with the quick convergence of the buyers' surplus share to nearly 100% in the NO Market, providing a first indication that obfuscation opportunities (i) severely weaken competition in the OO Market but (ii) do not completely remove it.

In addition, there is a large and stable spread in buyers' surplus from traded products throughout the whole 20 periods of the OO Market: the within-period standard deviation, measured as a share of the maximal total surplus, is 25.3% for offered products and 13.0% among traded products. Moreover, this dispersion does not differ much between the first 15 and the last 5 periods, while in the NO market dispersion quickly vanishes. This result shows again the much weaker competitive forces in the OO

---

[20] It appears that the main reason why the buyer surplus does not reach 100% is that sellers shy away from prices that give them literally zero profits, i.e., sellers want to earn at least 1 ECU from their trade.

[21] All our statistical results in the paper are based on t-tests that cluster standard errors at the market level unless specified otherwise.

Market and the fact that obfuscation opportunities alone can cause large and stable dispersion of buyers' surplus shares in traded products.

Figure 2: Average buyer surplus per trade as a share of the maximally possible surplus in NO Markets and OO Markets

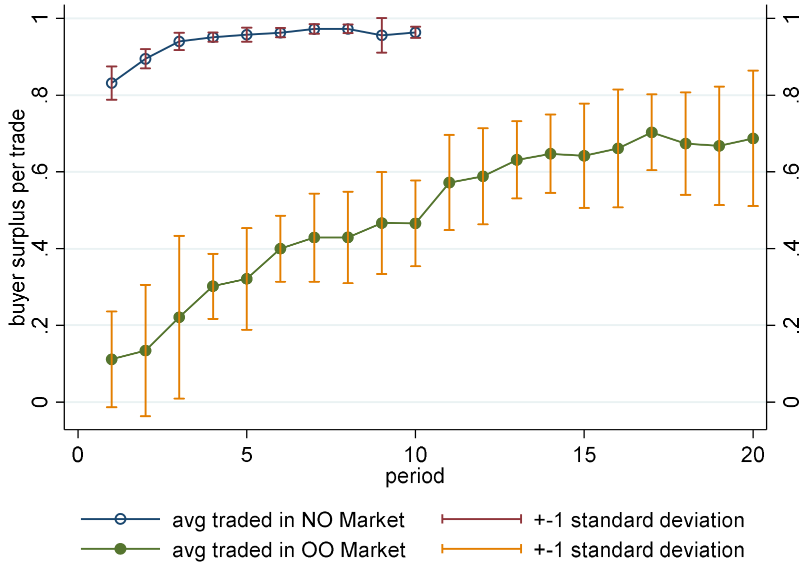

Notes: The figure shows the average buyer surplus per trade (plus/minus one standard deviation) as a percentage of the maximum possible total surplus across the 20 periods in the market without obfuscation (NO Market) and the market with obfuscation (OO Market). The figure is based on the pooled data from the Half-Cost Treatment and the Convex Cost Treatment. The traded buyer surplus in the OO Market is significantly lower than the buyer surplus in the NO Market ($p = 0.000$, t test).

Why do buyers only reap such a low share of the surplus in the market with obfuscation opportunities? Is it because sellers fail to provide the efficient level of extra features, or do sellers implement efficient extra features but charge high prices? Or is it because buyers do not identify the best offered product that is available in the market? The next result answers these questions.

Result 2 (Source of Buyers' Surplus Loss):

> In the market with obfuscation opportunities, buyers do not reap the whole surplus because (i) even the best available product in the market is associated with a substantial mark-up of prices over marginal cost, and (ii) buyers' persistent failure to purchase the best available product in the market.

Figure 3 shows the decomposition of the buyers' surplus loss in the OO Market described in Result 2. The figure displays (for each period) the average traded *buyer* surplus (the green circle line), together with the average *total* surplus from the best offer in the market (the blue triangle line) and the average *buyer* surplus provided by the best offer in the market (the red square line).[22] The best offer is defined by the product that gives the buyers the highest surplus. Thus, the triangle line shows the extent to which the best offer generates the maximal total surplus by implementing the efficient level of extra features. The difference between the total surplus and the buyer surplus in the best offered product in the market shows the share of the surplus that the sellers of the best offered products could appropriate. The

[22] All three graphs are percentages of the maximum possible total surplus; therefore, the scale of the vertical axis is the same as in Figure 2.

difference between the buyer surplus in the best offered product and the average traded buyer surplus in a period informs us about the extent to which buyers failed to buy the best offered product.

Figure 3: Decomposition of buyers' surplus loss in markets with surplus-enhancing obfuscation opportunities

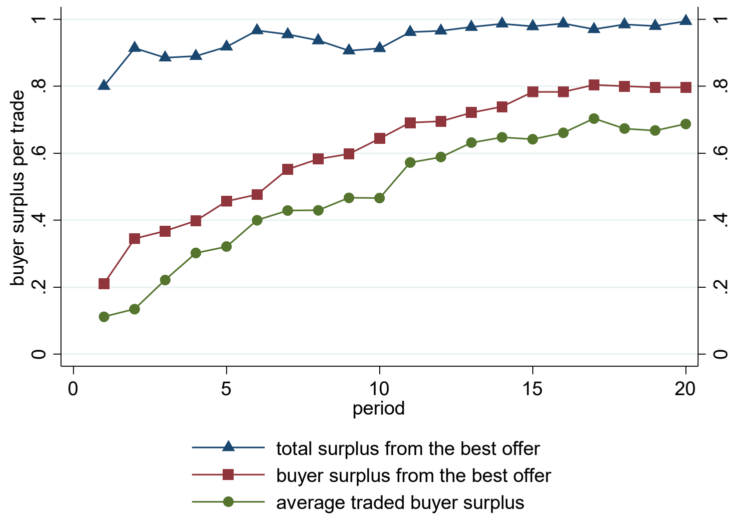

Notes: The figure shows the average buyer surplus in traded products, the surplus buyers could earn if they identify and buy the best offer in the market, and the total surplus generated by the best offer (product). The figure is based on the pooled data from the Half-Cost Treatment and the Convex Cost Treatment, and all graphs are displayed as a percentage of the maximum possible total surplus. The buyer surplus from the best offer is always lower than the total surplus from this offer (p = 0.000). Similarly, the average traded buyer surplus is significantly lower (p = 0.000) than the buyer surplus from the best offer.

Figure 3 illustrates the regularities described in Result 2. First, the big gap between the square line and the triangle line indicates that even sellers who make the best offer among the 6 competing products charge a price substantially above marginal cost and earn substantial profits. In Periods 16-20, the buyer surplus offered in the best available product stabilizes at only 80%, significantly lower than the maximal buyer surplus (p = 0.000). Second, the significant difference (p = 0.000) between the average traded buyer surplus (the circle line) and the buyer surplus in the best available product (square line) shows that, on average, buyers persistently fail to purchase the best offered product across the entire 20 periods of the OO Market. This suggests that it is persistently hard for buyers to identify the best product, i.e., in markets with obfuscation opportunities the buyers do not learn to make better purchase decisions over time. Consequently, this failure generates another loss in buyers' surplus of around 12%.[23] Overall, sellers in the OO Market appropriate reductions in buyer surplus from both high prices and buyers' failure to purchase the best offer, allowing sellers to earn around 32% of the total surplus in the long run (i.e., periods 16-20).

Why are sellers capable of appropriating a large share of the surplus and why do buyers on average fail to buy the best product in the market? One potential reason for this is that the OO Market is characterized

[23] The failure to identify the best offer has two sources. Buyers can fail to find the best offer (i) because they do not visit it and (ii) because they do not examine all add-ons of the visited products so that even though they visited the best product, they do not know that it is the best product. 54% of the surplus loss that is due to the failure to purchase the best offer is caused by reason (i) and in 46% of the cases buyers failed to buy the best product although they visited it but did not know that it is the best one (reason (ii)).

by complex products which make it difficult for buyers to compare products and find the best ones. Our next result indeed indicates that the market is characterized by a high degree of complexity:

Result 3 (Product Complexity):

(a) Product complexity is, on average, relatively high and persists throughout all periods. It converges to almost maximal complexity in the Half Cost Treatment (HCT) and to slightly less than 5 add-ons in the Convex Cost Treatment (CCT).

(b) In addition, sellers generally (i.e., in the HCT and the CCT) add more add-on features than needed to generate their desired level of feature values.

We provide support for Result 3 in Figures 4a and 4b. Figure 4a shows the average complexity of both offered and traded products in HCT and the CCT, respectively. The average number of offered add-on features per product is 5.9 in HCT and 4.7 in CCT in periods 11-20, and the complexity of average traded products is even slightly higher than those of the offered products although this difference is not significant. Recall that the efficient aggregate feature value in the HCT requires 6 add-ons, while efficiency only requires 4 add-ons in the CCT. However, the average offered product complexity is at a stable level significantly above 4 (p = 0.009) in the CCT, meaning that many sellers are willing to add too many features relative to the surplus-maximizing level of feature values.

Figure 4: Product complexity and over-complexity in the Half Cost Treatment (HCT) and the Convex Cost Treatment (CCT) of markets with surplus-enhancing obfuscation opportunities

(a) Product complexity **(b) Over-complexity**

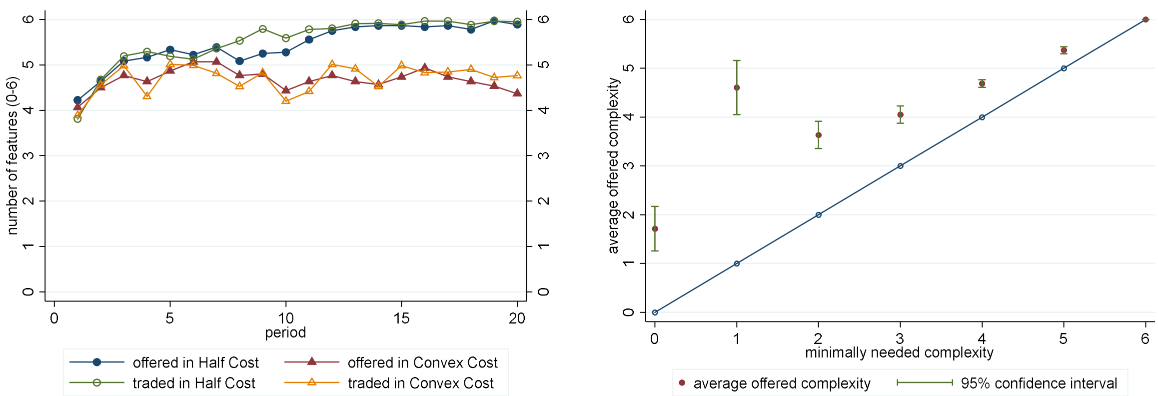

Notes: Figure 4(a) shows the average number of features in both the offered and traded products in the OO Market. The figure presents data from the Half-Cost Treatment (HCT) and the Convex Cost Treatment (CCT) separately. The average offered product complexity in CCT is at a stable level significantly above the surplus-maximizing level of 4 (p = 0.009); but it is significantly below the average offered product complexity in HCT (p = 0.006). Figure 4(b) shows the average complexity of offered products (together with the associated 95% confidence interval) compared to the minimal complexity needed to generate the offers' planned feature values. The figure is based on the pooled data from the Half-Cost Treatment and the Convex Cost Treatment.

Results 3(b) is illustrated in Figure 4(b). In this figure, the horizontal axis categorizes all the offered products according to how many features sellers minimally need to produce their desired aggregate feature value, while the vertical axis gives the actual average number of features that these products have. Thus, if sellers offer only products with the number of features that are minimally required to produce the desired feature values, the actual average number of features should be on the 45-degree line. In fact, however, all the offered products with a minimally required feature number of 5 or less are significantly more complicated than needed to produce the desired aggregate level of feature values.

Taken together, the market-level results indicate that sellers earn close to zero profits in the NO Markets whereas in OO Markets, where buyers have to search through products with many add-on features, no seller prices the product at marginal cost, and even inferior products have a fair chance of being sold because many buyers persistently fail to buy the product with the highest buyer surplus.

3.2 Buyers' search and purchase behavior

Why are sellers able to reap a substantial part of the total surplus in the OO Market, and why is the law of one price violated? A natural answer to this question is that the complexities generated by add-ons in the OO market impose "search" burdens on buyers because it takes time to find, understand, and evaluate the available products. Therefore, we examine the buyers' search and buying behavior in the OO Market next.

Result 4 (Buyers' Search Behavior):

(a) In a substantial number of cases, buyers in the OO Market visit[24] only one or two products in the market and study these products in depth but, overall, there is substantial heterogeneity in the number of products visited and the percentage of add-ons studied.

(b) Buyers' search behavior indicates a strong trade-off between "browsing" (the number products visited) and "studying" (the depth with which individual products are examined). More intense browsing is associated with a considerable reduction in studying.

(c) Buyers whose time needed for assessing the overall value of a product is more convex in the number of add-ons are relatively more focused on browsing. Buyers whose mental cost of assessing the overall value is higher are also relatively more focused on browsing.

Support for Result 4a is provided in Figure 5a and 5b. Figure 5a shows the distribution of the number of products visited in the OO Market; it indicates that in roughly 60% of the cases only 1 or two products are visited while in roughly 35% of the cases, the buyers visit between 3 and 6 products, which indicates substantial heterogeneity in search behavior. This result is thus in line with the fundamental idea in many theories (e.g., Ellison, 2005; Heidhues, Johnen and Köszegi, 2021) that assume or predict very limited search by the buyers (i.e., one or fewer visits) in response to product complexity. In this context, it is noteworthy that we observe that in 42% of the cases buyers indeed visit only one or fewer products. But our results also indicate many buyers make more than one visit.

In addition to heterogeneous breadths of search, buyers may also differ in the extent to which they "study" the visited products in depth. Figure 5b illustrates that in 50% of the cases in which a product is visited, buyers study between 80% and 100% of the add-on features. Thus, in-depth study of a limited number of products is a frequent behavior. However, like for the breadth of search, there is substantial heterogeneity in buyers' in-depth study of products as in roughly 50% of the cases, buyers study less than 80% of the add-ons (see Figure 5b), i.e., they obtain only a partial understanding of the product's add-on features. Both search patterns shown in Figure 5a and 5b do not change much across the 20 periods.

[24] A product is visited if the buyer examines at least one of the product's add-ons.

Figure 5: Buyers' search behavior

(a) Number of visited products

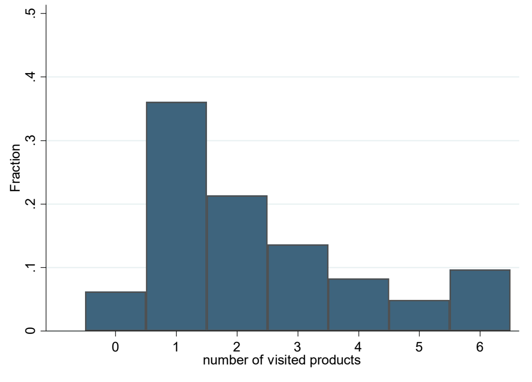

(b) Percentage of add-on features seen

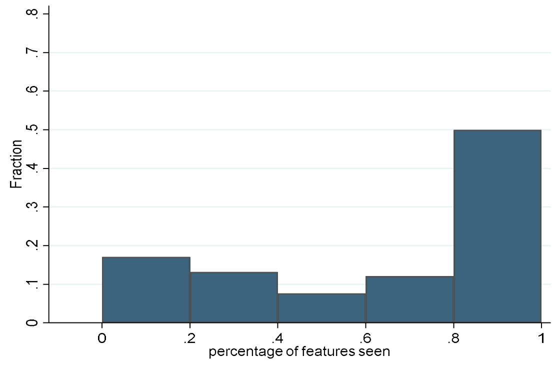

(c) Percentage of features seen conditional on visiting a given number of products

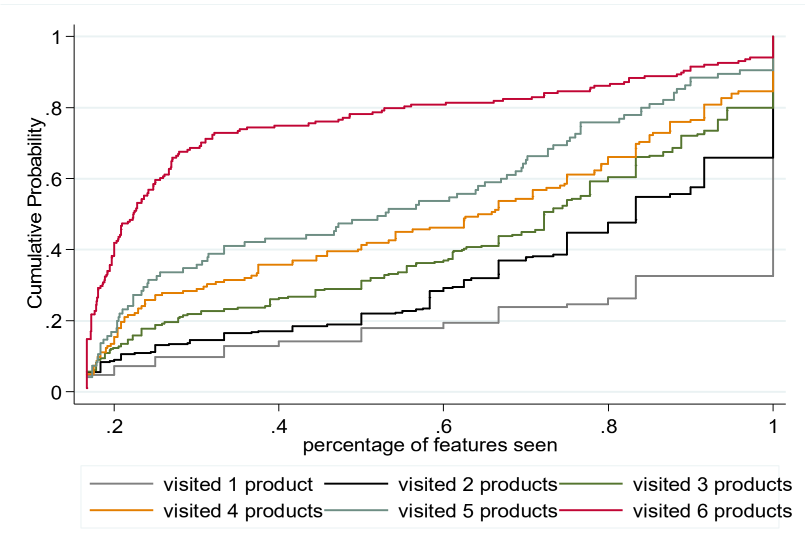

(d) Relative search focus and within-product convexity of time needed

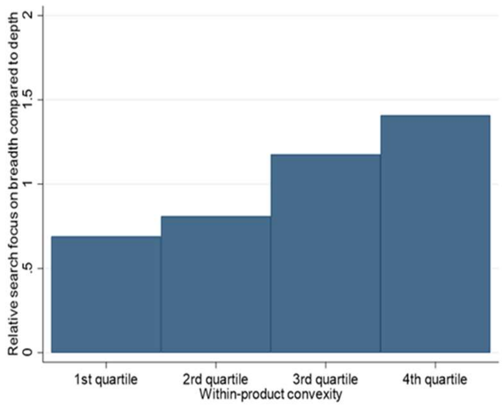

(e) Relative search focus and mental cost

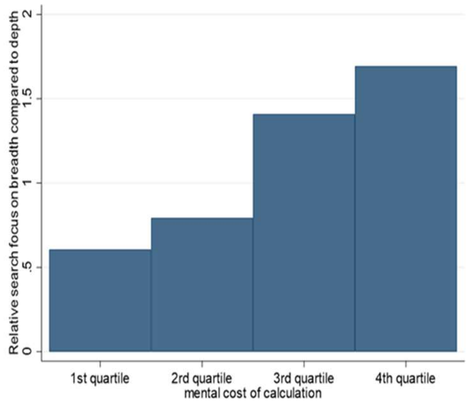

Notes: The figure shows the distributions of (a) the numbers of products that buyers visit, (b) the average percentage of features that buyers see for the products that they visit. Figure 5(c) shows the percentages of features that buyers see for those products that they visit, conditional on the numbers of products that they visit. Figure 5(d) displays buyers' relative search focus as a function of buyers' within-product convexity of the time needed to compute the overall value of a product. The relative search focus is measured as "percentage of products visited divided by the percentage of add-ons examined". Within-product convexity is measured by the "time needed per add-on for products with 6 add-ons divided by the time needed per add-on for products with 2 add-ons". Buyers are partitioned in 4 quartiles according to their within-product convexity. Figure 5(e) displays buyers' relative search focus as a function of buyers' mental cost of assessing the overall value of a product.

Support for Result 4b is provided by Figures 5c. This figure displays the cumulative distribution of add-on features that are examined, conditional on visiting 1, 2, 3, etc. products. It documents the existence of a strong tradeoff between the depth of search (studying) and the breadth of search (browsing). For example, conditional on visiting only one product the depth of search is very high, i.e., between 80-100% of the add-ons are examined in roughly 70% of the cases. In contrast, if buyers visited all 6 products their depth of search is very low because in roughly 75% of the cases, they examined less than 40% of the add-ons. More generally, Figure 5c shows that the cumulative probability of examining only a small percentage of add-on features is monotonically increasing in the number of products visited.

The existence of a tradeoff between studying and browsing rationalizes an important assumption in the model by Heidhues, Johnen and Köszegi (2021) in which the lack of browsing is an important competition-limiting force. We find that this trade-off is primarily an aggregate-level phenomenon because at the individual level, given the number of products that subjects visit, they examine on average all their visited products in similar depths, i.e., the trade-off between browsing and studying is not a within-subject phenomenon. The trade-off is rather due to different types of buyers. A majority of buyers focus their search predominantly on depth: they visit a small number of products that are examined in depth. In contrast, a substantial minority of buyers examine products in a relatively superficial way but visit many products.

But how can we make sense of heterogeneous buyer types that display a systematically different relative search focus? Is it possible to provide a deeper rationale for their existence? To examine this question, we analyzed data from Part 3 of the experiment that allows us to elicit subjects' *per add-on* time needed for assessing the overall values for products with varying numbers of add-ons. For a product with more add-ons, the computation is presumably more cognitively demanding or effortful. Thus, subjects may differ in the extent to which the time needed to compute the overall value of a product per add-on varies with the number of add-ons. We find indeed (see Figure 5d) that subjects with a higher convexity of time needs display a stronger relative focus on browsing.[25] Note that this makes perfect economic sense because for those with a higher convexity it is relatively more time-consuming to examine a smaller number of products more deeply instead of examining more products superficially.

Thus, despite a constant monetary unit cost of time these subjects have an economic reason to focus on browsing relative to studying. We also find (Figure 5e) that buyers' purely mental cost of assessing the overall value of products is also positively correlated with buyers' relative focus on browsing.

Taken together, Result 4 indicates that buyers remain very incompletely informed about the prevailing products in the market – either because they visit only a few products or because they view only a few features of the visited products. This means that buyers cannot respond to changes in add-on features of products they have either not visited or not examined in depth. As a consequence, one would expect that the demand for products responds sluggishly to the values and prices of a product's add-ons as well as to the values and prices of the add-ons of a given product's competitors. In contrast, because base prices

[25] The presence of convex time needs is reminiscent of the findings in Altmann et al. (2022) who show the existence of negative spillovers between a memory task (remembering a 7-digit number) and a subsequent math task. In our set-up, where in-depth evaluation of a product with many add-ons may be viewed as a series of math tasks (adding the net values of add-ons) that are preceded by a memory task (memorizing the net value of each previous add-on, or memorizing the preliminary overall net value of all previous add-ons), the convexity of the time needs may be interpreted as a negative spillover effect of the previous task on the subsequent one.

are transparently visible in the market, product demand is expected to be more responsive to base prices. Our next result documents whether buyers' purchase behavior indeed confirms these expectations.

Result 5 (Buyers' Demand Behavior):

> **(a)** Buyers' demand for a product substantially increases if the product's base price decreases. Likewise, the base price of the best competing product in the market also has a sizeable impact on a product's sales.
>
> **(b)** The aggregate values and prices of a product's add-on features exert a considerably smaller effect on the product's sales compared to base prices. In fact, product demand is three-four times more responsive to base prices than to aggregate feature prices. Likewise, the impact of the aggregate value and aggregate price of the add-ons of the best competing product on a product's sales is relatively small.

Evidence for Result 5a and 5b is provided in Table 2 which reports an OLS regression of product units sold on the characteristics of the product itself and the characteristics of the best available competing product in the market. The regression also control for market complexity (i.e., the average number of extra features) and time trend; standard errors are clustered on the market level. It turns out that base price variations exert the largest impact on a product's sales: lowering the base price by 10 ECUs significantly increases sales by 1.8 units. Similarly, the base price of the best available (best competitior's) product in the market plays a similar (although slightly smaller) role: a reduction in the best competitor's base price significantly lowers a seller's sales.

The high elasticity of product demand to headline prices contrasts with their much lower elasticity to the aggregate feature values and aggregate feature prices: an increase in aggregate feature prices by 10 ECUs reduces a product's sales only by 0.6 units, indicating that product demand is three times more responsive to base prices. Likewise, an increase in aggregate feature values by 10 ECUs increases a product's sales only by 0.7 units.[26] The difference in the response to headline prices versus the response to aggregate features values and prices is highly significant ($p = 0.000$).

Furthermore, the aggregate feature values and prices of the best competing product in the market have also small impacts on product demand; and the best aggregate feature value or the best aggregate feature price among the competitors basically does not matter at all. It is worthwhile to contrast these sluggish demand responses to competitors' aggregate feature values and prices with our findings in the NO Market. The best available product in that market attracts almost all the buyers, practically eliminating the sales of competing sellers, while competition is strongly mitigated and inferior products have a good chance of being sold in markets with obfuscation opportunities.

Before we move on to the next section, we ask whether buyers' purchase behavior is not only affected by the base prices and aggregate feature values and prices, which have direct consequences for the buyers' monetary payoffs, but whether buyers also exhibit a preference for the number of add-ons per se or specifically labelled add-ons. Notice that Table 2 shows that product demand is not affected by the number of a product's add-ons, indicating that buyers do not exhibit a significant preference for the number of add-ons per se. In addition, we show in Appendix 3 that the sellers also believed that buyers

[26] Notice that unlike in Chetty et al. (2009), in our set-up this low responsiveness of buyers to aggregate feature values and prices cannot be due to their unawareness of the *existence* of "hidden" values and surcharges. Instead, this low responsiveness prevails even though buyers know from the experimental instructions that add-ons may exist. The low responsiveness is therefore a result of buyers' limited breadth and depth of search.

do not have a significant preference for the number of add-ons but that they predominantly care for their monetary payoff. We also performed extensive analyses and use survey evidence from questionnaires in Appendix 3 to examine the influence of specific feature labels on buyers' product demand. The results in this appendix show that the role of labels is negligible for buyers' behavior. In addition, sellers rationally anticipate this such that the role of labels is also negligible in sellers' product design decisions.

Table 2: Buyers' responses to a product's own and the competing products' characteristics

Dependent Variable	Units sold
Base price	-0.18*** (0.01)
Aggregate feature value	0.07*** (0.01)
Aggregate feature price	-0.06** (0.01)
Best competitor's base price	0.12*** (0.01)
Best competitor's aggregate feature value	-0.08*** (0.01)
Best competitor's aggregate feature price	0.07*** (0.01)
Number of features	-0.05 (0.07)
Average number of features among competitors	-0.12 (0.09)
Period	-0.03** (0.01)
Constant	3.70*** (0.33)
No. of observations	1320
R-square	0.23

Notes: Standard errors are clustered on the market level. *** indicates significance at the 0.01 level, ** at the 0.05 level, and * at the 0.10 level.

3.3 How does buyers' search and purchase behavior affect sellers' incentives and behavior?

To assess the incentives sellers faced when determining the base prices, the aggregate feature prices and values as well as the ordering of the attractiveness of their products' individual add-ons, it is important to know whether sellers anticipated that buyers' depth and breadth of search is rather limited. To examine the sellers' beliefs about buyers' search behavior we asked the sellers several questions in Part 3 of the experiment.

With regard to the breadth of search we asked the sellers the following question: "*In every period, buyers could visit at most six products by looking at some of their extra features. On how many of these six products, do you believe, did they take a look?*". On average, sellers believed that buyers examined 3.23 of the 6 products, indicating a strong belief in the limited breadth of buyers search behavior.

With regard to the depth of search we asked sellers: "*You are asked to guess for those products that buyers visited, how many features they, on average, examined before they visited another product or* *made a purchase. For products with 2 features, how many features did buyers on average examine? For products with 4 features, how many features did buyers on average examine? For products with 6 features, how many features did buyers on average examine?*" Sellers believed that buyers examine on average 1.88 features of products with 2 add-ons; 2.93 features of products with 4 add-ons and 3.26 features of products with 6 add-ons. Note that this result shows that sellers' believethat products with a higher number of add-ons will generate a larger share of non-examined add-ons.Taken together this means that sellers were well aware that buyers' depth of search was, on average, rather limited, and that the lack of relative depth increases with the number of add-ons.

The fact that sellers anticipate the limited depth and breadth of buyers' search behavior has important implications for the incentives the sellers face when choosing the number of add-ons and their values and prices. In particular, sellers' belief that many buyers fail to examine all add-on features creates incentives to manipulate the ordering of the add-on features by placing the best add-ons first and the worst add-ons last. In this way, sellers can increase the probability that buyers do not notice the worse add-on features. Recall that the sellers can achieve this by re-randomizing the order of add-on display.[27] We will examine whether sellers behave in line with this incentive in the next result.

Moreover, because sellers believe that products with a higher number of add-ons will be associated with a larger share of non-examined add-ons, they have an incentive to increase product complexity beyond what is needed to generate the products' intended values. We have already shown in Result 3a and Figures 4a and 4b that sellers' complexity choices are indeed consistent with these incentives.

Finally, sellers' beliefs also imply that they anticipate that competition with regard to aggregate feature prices and values is relatively muted because many buyers are not even aware of competing products' add-on features. In contrast, sellers know that buyers can easily access base prices such that buyers' responses to base prices will be more elastic and, thus, competition in base prices will be stronger.

Many theoretical models (e.g., Ellison 2005; Spiegler, 2016; Choi, Dai and Kim 2017; etc.) have conjectured that headline prices may serve as an attention-grabber in an add-on pricing setting and may, therefore, be set at very low levels in order to attract consumers. Once consumers have been lured to visit a firm's store or website, they may have a limited willingness to switch to competing firms, which is exactly the behavior displayed by many of our subjects (the "studiers"). However, as the attractiveness of base prices and "hidden" product features could be negatively correlated (Célérier and Vallée 2017; Shulman and Geng, 2019), buyers may also interpret low base prices as a signal for highly priced "hidden" product features, which may weaken the overall attractiveness of very low base prices. Our next result documents how headline prices and add-on prices are determined by sellers.

Result 6 (Ordering of Add-ons and Determination of Base and Aggregate Feature Prices):

(a) Sellers order the add-on features in the OO Markets so that those features that first become visible are most attractive for the buyers, while the least attractive add-on features only become visible through deeper search.

(b) Initially, the average base prices are considerably above marginal cost, but they gradually fall over time and eventually they stabilize below the base products' marginal cost. However, base prices remain dispersed throughout the 20 market periods.

[27] Recall that we simplified the sellers' choices such that they only determine the aggregate feature value and the aggregate feature price of their product, while the computer determines the values and prices of individual features randomly. If sellers are not satisfied with one realization, they can re-randomize until they are satisfied.

(c) Sellers are able to enforce aggregate feature prices substantially above costs, and the average profits sellers earn from add-on features are stable over time. At the same time, sellers do not appropriate all the surplus from add-on features, and net feature values remain dispersed across products.

(d) Low headline prices are on average associated with worse add-ons but the association is quite noisy.

Evidence for Result 6a is provided in Figure 6 below. If each of the individual feature values and prices were really randomly determined across the display positions of add-on features, then there should be on average no difference in the add-ons' net feature values v_f^i - p_f^i across display positions. That is, the average net value of the features shown at the top of the feature list should be the same as the average net value of the next shown feature, and so on.

Figure 6: Net feature values across display positions of add-on features

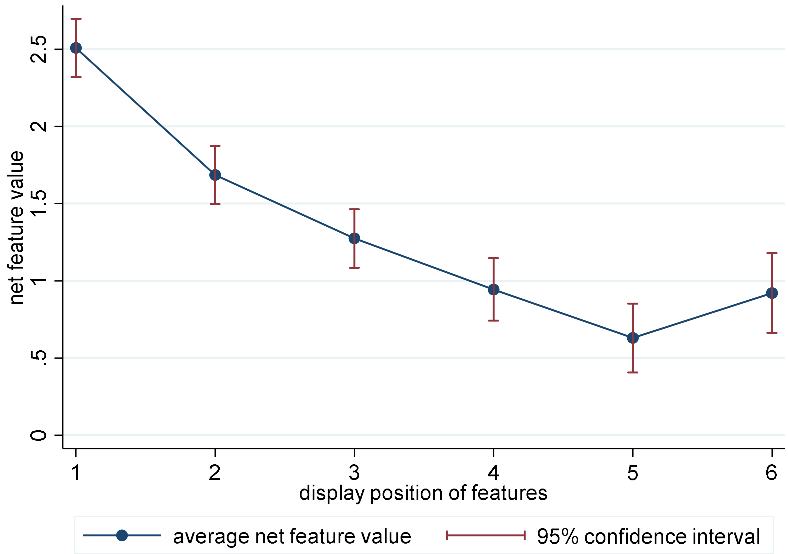

Notes: The figure shows the average net feature value (i.e., feature value – feature price) of individual add-ons in offered products across display positions of individual features within a product. The associated 95% confidence intervals are also presented. The figure is based on the pooled data from the Half-Cost Treatment and the Convex Cost Treatment. A joint t-test for whether the net feature value of each feature is higher than its next feature yields a p-value of 0.001. The first 4 features are also individually significantly (at least at a 10% level) different from each other (p = 0.001 between the first and the second feature, p = 0.072 between the second and the third feature, and p = 0.016 between the third and the fourth feature).

However, Figure 6 shows that this is not the case. The figure depicts the average net value of features on the y-axis across display positions of features on the x-axis. The figure shows that at any position of the display order, that position's feature is on average significantly more attractive than the next position's feature[28] (p = 0.001 from a joint clustered t test between features in one position and in the next position). This pattern is particularly pronounced among the first three feature positions. Thus,

[28] Only the last 2 features are not significantly different from each other; but this may simply result from the fact that, by design, sellers can determine the display order only by letting the computer re-randomize the whole order so that they do not have full control over the attractiveness of each extra feature.

because many buyers do not inspect all features of a product, they will only see the first few attractive features and remain uninformed about the product's less attractive features.[29]

Result 6b and 6c imply that headline prices are initially a source of profits but over time sellers eventually incur losses on their base products, while the add-on features become the key source of profits. Result 6b and 6c thus support models of add-on pricing that predict loss-leadership price structures.

Support for Result 6b comes from Figure 7a, which shows the average base prices in the OO Markets (jointly with the flat marginal cost curve and the buyers' valuations for the basic product). The figure illustrates that base prices are initially far above marginal cost but drop gradually and eventually fall somewhat below marginal costs from period 11 onwards (p = 0.616 and p = 0.020 for offered and traded base prices respectively). In fact, in the long run base prices stabilize at a level slightly below marginal cost. This is indicated by Figure 7a and the fact that neither the offered (p = 0.659) nor the traded base prices (p = 0.424) differ between periods 11-15 and 16-20. In addition, base prices remain rather dispersed even in periods 16-20. During these periods, the within-period standard deviation, measured as a percentage of the maximal total surplus, is 16.7% for the offered products and 12.1% for the traded products.

Figure 7: Base prices, aggregate feature values, prices and costs

(a) Base Prices

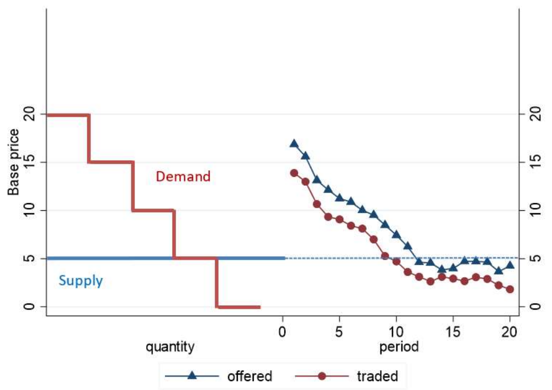

(b) Aggregate feature values, price and costs

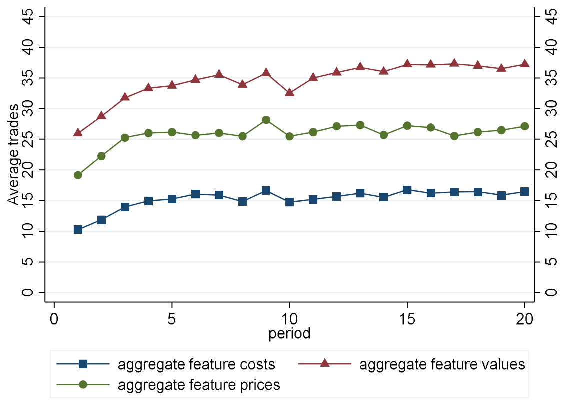

Notes: Figure 7a shows the constant marginal cost of the basic product (blue thick line in the left part of the figure), the buyers' valuations of the basic product (red thick line in the left part), and how the prices of the offered and traded basic products evolve over time (in the right part of Figure 7a). From Period 11 onwards, the average traded base prices are significantly lower than the marginal cost (p = 0.020). Figure 7b shows the aggregate feature values and the associated aggregate feature costs in traded products. The difference between aggregate feature values and aggregate feature costs is a measure of the surplus the extra features provide. The aggregate feature price (line with circles) shows how this surplus is distributed between buyers and sellers. Both figures are based on the pooled data from the Half Cost Treatment and the Convex Cost Treatment. Figure 7b documents that sellers appropriate a stable share of the surplus generated by the products' extra features. The aggregate feature price of traded products is significantly higher than their aggregate feature cost (p = 0.000). The aggregate feature value of is significantly higher than the aggregate feature price (p = 0.000).

Support for Results 6c comes from Figure 7b that illustrates how the average aggregate feature values, prices, and costs in traded products evolve over time. The difference between aggregate feature values

[29] This finding is congruent with field observations. Recall, e.g., the finding of Célérier and Vallée (2017) who document that financial firms hide the risks involved in their financial products behind complex product descriptions.

and costs reflects the fact that sellers generated a substantial surplus via add-ons. In addition, Figure 7b shows that aggregate feature prices turn out to be substantially and significantly higher than the aggregate feature costs (p = 0.000). Recall that the same technology to produce add-ons is available to all the sellers. Therefore, the extra features are highly replicable and, if competition is fully at work, they should not allow for prices above their marginal costs. However, sellers are able to obtain a sizable and stable share of profits from their add-ons over the entire 20 periods. Notice that this also means that the increase of the buyer surplus over time in the OO Markets (see Figure 2 and 3) stems entirely from declining base prices.

Result 6a – 6c provide strong support for the view that sellers respond quite rationally to the incentives generated by the buyers' search and demand behavior. Sellers' knowledge of buyers' highly incomplete depth of search generates incentives to show the worse add-on features last and to make products overly complex, which is exactly what we observe in our data. The highly elastic demand response to baseline prices induces sellers to compete strongly with their base prices while the less elastic demand response to aggregate feature prices and values strongly mitigates competition via add-on features.

However, although sellers appropriate a sizable share of the surplus from add-ons, they do not appropriate all the surplus. This makes sense in view of buyers' demand behavior stated in Result 5b which shows that higher add-on prices do generate a reduction in product demand, i.e., some competitive forces still affect the pricing of add-on features. Moreover, the within-period dispersion of aggregate net feature values across products is rather large: even in the long run (in period 16-20), the average standard deviation of the buyers' surplus from the add-on features is 20.0% of the maximal total surplus of traded products. This pattern of add-on feature prices differs from theoretical models that rely on the assumption of no competition among shrouded surcharges (e.g., Gabaix and Laibson, 2006; Heidhues, Kőszegi and Murooka, 2017; Heidhues, Johnen and Köszegi, 2021) and models that predict non-dispersed add-on prices that extract all the surplus (e.g., Lal and Matutes, 1994; Ellison, 2005). The large and persistent dispersion in aggregate net feature values could have important implications for how markets with add-ons function. First, the existence of dispersed net feature values may foster competition with regard to add-on prices as buyers may find better products if they search more. Second, dispersed net feature values also may render base prices – which constitute only a small part of the product's overall properties – an imperfect signal for the total net feature value of the products.

Although both base prices and net values of add-on features are quite dispersed in the market, there may still be a relationship between their average levels, based on which buyers may infer the aggregate value of add-on features from the product's base prices. We previously discussed the possibility that low headline prices may indicate that sellers provide overpriced "hidden" add-ons. If this were the case, buyers might become suspicious of products with low headline prices, which would constrain sellers' ability to lure buyers into buying products with overpriced add-ons. Result 6d shows that low headline prices are indeed associated with a lower overall net value of add-ons for the buyers: In a regression of headline prices on products' aggregate net feature values (i.e., aggregate feature values – aggregate feature prices) that controls for period dummies to eliminate the time trend, the regression coefficient is positive and significant (p = 0.024). However, the regression coefficient is not very large (0.43) and there is considerable noise in this relationship which may make it difficult for buyers to make strong inferences (see Figure A4 in the Appendix 4). Nevertheless, some buyers may have had bad experiences when buying a product with a low base price that has unattractive add-ons which could well have induced suspicion among them.[30] This suspicion, in turn, may explain why average base prices stopped falling after period 10.

3.4 How do buyers and sellers respond to an increase in search cost?

Taken together, our results suggest that buyers and sellers respond quite rationally to the prevailing economic incentives. Facing a market with rather complex products, the buyers adapt their search focus to their pecuniary (i.e., time) and their mental cost of search.

Buyers whose time needed to study additional add-ons of a given product is a highly convex function of the number of add-ons, and buyers with high mental cost of search tend to be "browsers" who visit many products but examine them only in a superficial way. In contrast, those with little convexity in the time needed for examining additional add-ons and with low mental cost of search tend to be "studiers"; they examine a few products in full depth.

Sellers rationally anticipate that buyers' breadth and depth of search is rather incomplete which provides – in addition to the surplus-enhancement through complex products – incentives to keep product complexity high, and to show the most attractive add-ons first. These are exactly the seller behaviors we observe in the OO-Market. In addition, seller competition is relatively strong in base prices where buyers' responses are quite elastic, but relatively weak in aggregate feature prices where buyer responses are much less elastic. Taken together this enables the sellers to enforce aggregate feature prices that are substantially above marginal cost.

To provide further evidence for our interpretation that buyers and sellers responded rationally to the incentives they faced, we conducted the "robustness check treatments" (see Table 1) where participants were randomly allocated to either a low search cost treatment or a high search cost treatment in the CCT. In the high-cost treatment, the buyers faced twice the time cost of search compared to the low-cost treatment. If buyers and sellers indeed respond rationally to the economic incentives this increase in the pecuniary search cost should generate very specific responses. Compared to the low-cost treatment, buyers in the high-cost treat should, in particular,

(i) reduce the overall time spent searching in the market, which should

(ii) reduce the breadth and the depth of search, and thus increase buyers' information imperfections. This, in turn, should further

(iii) reduce buyers' responsiveness to the changes in add-on prices and values of a given product as well as their responses to the add-on prices and values of the best competing product.

Finally, when facing buyers who search less as predicted in (i) – (iii),

(iv) sellers should enforce higher aggregate feature prices for their add-ons such that they reap a higher share of the total surplus from the add-ons.

[30] Further evidence on buyer suspicion against very low base prices is provided in Appendix 4 where we document that the average base price of the first visited products as well as the average base price of the traded products is significantly higher than the lowest base price in the market. This evidence suggests that a non-negligible number of buyers are not attracted by the lowest base price.

We find that all four predictions are strongly supported by the data. Below we summarize these behavioral changes and in Appendix 5, we graphically illustrate the changes:

Result 7 (Behavioral and Market Responses to a Search Cost Increase):

(a) In each period, the time spent searching in the high-cost treatment is lower compared to the low-cost treatment (see Figure A7, and the search cost increase significantly ($p=0.004$) reduces the average time spent searching in the market by 34%.

(b) The search cost increase shifts the whole distribution of individuals' depth of search to the left (see Figure A.8) and, on average, the depth of search is reduced by 20.9% ($p < 0.036$). Likewise, the distribution of the breadth of search shifts to the left (see Figure A.9) and the average breadth is reduced by 13.5% ($p < 0.020$).

(c) Higher search costs generate a large reduction in buyers' responsiveness to add-on prices (see Table A.2 in Appendix 5). In particular, at higher search cost the buyers' response to a rise in a product's aggregate feature price (aggregate feature value) is reduced by 59% (56%). Likewise, the buyers' response to the aggregate feature price (aggregate feature value) of the product's best competitor is even reduced by 71% (64%). All these reductions in buyers' responses to price changes are significant (all $p < 0.03$).

(d) Higher search costs significantly increase the sellers' share of the traded surplus from add-ons from 35.3% to 58.9% ($p = 0.057$, see Figure A10), indicating that sellers are able to enforce a considerably higher mark-up over aggregate feature cost.

4 The impact of surplus-neutral add-ons

In the Surplus-Neutral Treatment (SNT), the experiment approximates a situation typically modelled in the theoretical literature on obfuscation. In this case, adding additional product features only raise complexity but do not come with enhanced surplus, and the decision to make the product complex or not is independent of the overall pecuniary value and price of the product. While the data of our treatments with surplus-enhancing extra features already showed many implications of an obfuscated market and supported many qualitative predictions of these theories, we now examine whether our key results on obfuscation remain robust when obfuscation does not add surplus. We summarize our main findings in the following.

Result 8: (a) Product complexity is significantly lower with surplus-neutral extra features compared to markets with surplus-enhancing extra features. In addition, high levels of market complexity are more fragile in the Surplus-Neutral Treatment in the sense that complexity decreases over time to low levels in some markets, eventually leading to competitive market outcomes.

(b) In these low complexity markets, buyers display an aversion against complex products which appears to induce sellers to offer products with lower complexity.

(c) A higher number of extra features is associated and believed (by buyers) to be associated with *better* products in markets with surplus-enhancing add-ons. In contrast, in markets with surplus-neutral add-ons, a higher number of add-ons is associated and believed to be associated with *worse* products. This treatment difference may explain why obfuscation is more fragile in markets with surplus-neutral add-ons.

Support for this result comes from Figure 8 below displaying the complexity levels and the fact that we observe a bi-modal complexity pattern across the 6 SNT markets: in 3 of the 6 markets (henceforth labelled "high complexity markets"), the average number of extra features across the entire 20 periods is 5.21, 4.63, and 4.73, respectively, while the average complexity levels in the other 3 markets (henceforth labelled "low complexity markets"), are only 3.18, 2.79, and 2.05, respectively. Moreover, the number of add-ons in the high complexity markets increases slightly over time (see Figure 8), while it declines in the low complexity markets, although the initially offered complexity levels in the first three periods do not differ significantly in the two types of markets ($p = 0.114$). Consequently, the average number of extra features in offered and traded products is higher in the high complexity compared to the low complexity markets ($p = 0.001$ for offered products, $p = 0.006$ for traded products). Due to the existence of the low complexity markets in the Surplus-Neutral treatment, the overall average number of features in offered and traded products in the SNT is lower than in the treatments with surplus-enhancing extra features ($p = 0.020$ for offered and $p = 0.025$ for traded products).

Several notable patterns arise in the SNT. First, in high complexity markets, the average complexity level is very similar to that in the treatments with surplus-enhancing extra features ($p = 0.304$). Consequently, sellers in these high complexity markets are able to appropriate similar levels of buyers' surplus as in HCT and CCT (see Figure A12 in the Appendix 8). Second, Figure 8 also shows that the complexity levels in offered products and actually traded products are quite similar in high complexity markets of the SNT, a pattern that is also present in HCT and CCT. In contrast, in the low complexity markets buyers on average buy products that are considerably less complex than the average product offered in the market, a discrepancy that persists across the 20 periods (see Figure 8). In the long run (periods 16-20), the average complexity of offered products converges to only 2.3 extra features and the buyers buy products with only 1.4 extra features in the low complexity markets. Due to the high transparency in these markets, market outcomes eventually become quite competitive, and buyers appropriate almost all the surplus in the market (see Figure A12). Taken together, Figure 8 and Figure A12, along with the associated statistical tests, provide support for Result 8a.

Why are sellers unable to sustain a high level of complexity in the low complexity markets? Figure 8 shows that they tried to increase product complexity during the first few periods as in the high complexity markets. Moreover, the sellers in the low complexity markets – like those in the high complexity markets – also strategically place the best add-on features on the top of the add-on list while the bad add-on features are "hidden" by placing them on the bottom of the list.[31] These findings suggest that sellers in both markets attempted to take advantage from buyers' limited depth of search by hiding the worse add-on features.

However, Figure 8 hints at a potential explanation for the unraveling of complexity levels in the low complexity markets. It shows that buyers in the low complexity markets tend to buy products with much lower complexity than those offered, a pattern that is absent in all other OO Markets. To explore this behavior more carefully, we examine the determinants of buying behavior econometrically analogously to Table 2. To characterize the extent to which buyers in the low complexity markets behave differently than buyers in the high complexity markets, we interact the determinants of buying behavior studied in Table 2 with a dummy variable that takes on the value of 1 if the observation comes from the low complexity markets. Our results are displayed in Table 3.

[31] Specifically, the net value from add-on features declines on average by 0.52 ECU for each consecutively displayed feature in the low complexity markets. In the high complexity markets this number is 0.39 ECU.

Figure 8: Number of add-ons in the market with surplus-neutral add-ons

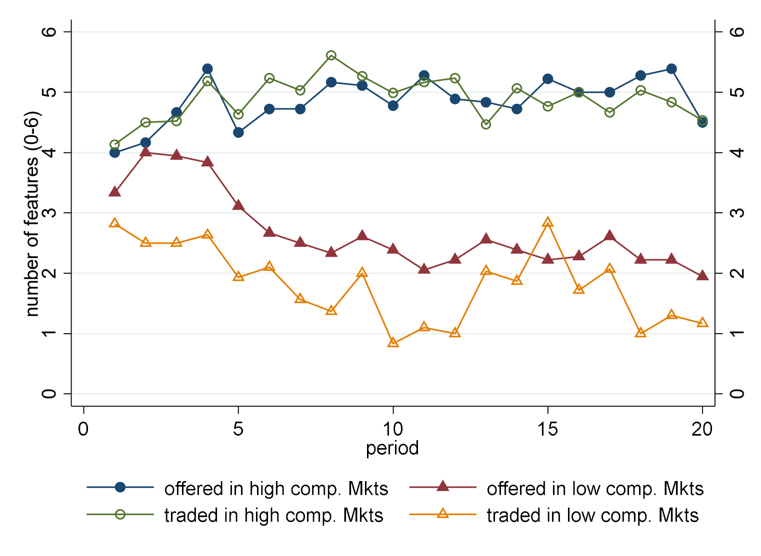

Notes: The figure shows the average number of features in both the offered and traded products in the OO Markets of the Surplus-Neutral Treatment. The average number of extra features in offered and traded products is higher in the high complexity compared to the low complexity markets (p = 0.001 for offered products, p = 0.000 for traded products).

Table 3 replicates important insights we already observed in Table 2 (which reports the same regression with the pooled data from HCT and CCT). In particular, the products' own base price and the base price of the best competing product are highly influential and significant determinants of product demand in both the high and the low complexity markets of the SNT. However, the table also highlights key differences between the high and the low complexity markets. Most importantly, while the coefficient on the product's "number of features" is insignificant in the high complexity markets, the number of features has a large negative effect on a product's sales in the low complexity markets. Controlling for all other characteristics of a product (such as base price, aggregate feature price, aggregate feature value) and for the characteristics of the best competitor's products, the addition of two more extra features to a product reduces product demand by 1.2 on average. We interpret this fact as an indication of buyers' aversion against complex products because it shows up even though we control for all other characteristics of a product and the characteristics of the best competing product.[32] Moreover, if it is indeed the case that buyers' in the low complexity market dislike complex goods, then we should also observe that the average number of features among the competitors raises the demand for a seller's own product, which is exactly what we observe: the coefficient on the interaction term between the average number of features in competitors' products and the "low complexity market" dummy is positive, large, and significant. Finally, complexity aversion already seems to be present in the low complexity markets from the very beginning because Figure 8 indicates that the complexity of the traded products is already lower than the complexity of the average product in the market during the first few periods. Taken together, these patterns support Result 8b.

[32] Therefore, like in the treatments with surplus-enhancing add-on features, there is no evidence that buyers prefer more add-on features per se – a conclusion that is further corroborated by the evidence presented in Appendix 3. Thus, in none of our treatments sellers' decisions regarding add-on features can be explained by buyers' non-pecuniary preferences for products with more features.

Table 3: Buyers' responses to a product's own and the competing products' characteristics in SNT

Dependent Variable	Units sold
Base price	-0.12** (0.04)
Base price × low comp. mkts	-0.14** (0.05)
Aggregate feature value	0.05 (0.07)
Aggregate feature value × low comp. mkts	0.16* (0.07)
Aggregate feature price	-0.07 (0.04)
Aggregate feature price × low comp. mkts	-0.10* (0.05)
Best competitor's base price	0.11** (0.03)
Best competitor's base price × low comp. mkts	-0.03 (0.07)
Best competitor's aggregate feature value	-0.07* (0.04)
Best competitor's aggregate feature value × low comp. mkts	0.04 (0.07)
Best competitor's aggregate feature price	0.08 (0.04)
Best competitor's aggregate feature price × low comp. mkts	-0.06 (0.08)
Number of features	0.13 (0.15)
Number of features × low comp. mkts	-0.58** (0.21)
Average number of features among competitors	-0.15 (0.13)
Average number of features among competitors × low comp. mkts	0.46** (0.16)
Period and Period × ow comp. mkts	√
Constant and Constant × how comp. mkts	√
No. of observations	720
R-square	0.29

Notes: The standard errors in the regressions are clustered on the market level. *** indicates significance at the 0.01 level, ** at the 0.05 level, and * at the 0.10 level.

One further noteworthy aspect in Table 3 is the following: Recall that in the markets with surplus-enhancing extra features (i.e., in the HCT and the CCT), a rise (decline) in the products' own aggregate feature value (aggregate feature price) by 10 units increases a product's sales by 0.7 (0.6) units. The corresponding increase in the product's sales is very similar at 0.5 (0.7) units in the high complexity markets.[33] This confirms that the elasticity of demand with regard to add-on characteristics is relatively low in markets with high complexity levels (like the HCT, the CCT and the high complexity markets of the SNT). In contrast, the market quickly becomes much more transparent in the low complexity markets of the SNT because of the lower complexity level, and this higher transparency may render a product's

[33] These coefficients are not significant in the high complexity market of the SNT, but this is likely to be due to the fact that we only have three of these markets and we cluster standard errors on the market level, while we have six markets in both the HCT and the CCT.

sales more elastic to the (smaller number of) add-on characteristics. The remarkably large interaction terms between the "low complexity dummy" and the products aggregate feature values or aggregate feature prices is in line with this conjecture. For example, while a rise in the products' aggregate price by 10 units reduces product demand by 0.7 units in the high complexity markets, the corresponding reduction is 1.7 units in the low complexity markets.

The previous analyses suggest that buyers' complexity aversion is likely to be a reason for the fragility of high complexity levels in the Surplus-Neutral Treatment (SNT). But why are buyers averse to complex products in the SNT, but not in the treatments with surplus-enhancing add-ons? Figure 9 below provides a potential answer to this question. Buyers in the treatments with surplus-enhancing add-ons experience and believe that there is a positive relationship between the products' aggregate net feature values and the number of add-ons (Figure 9a). In contrast, buyers in the SNT experience and believe on average that this relationship is negative (Figure 9b). Thus, a higher number of features is associated and believed to be associated with worse products in the Surplus-Neutral Treatment[34], while the opposite is the case in the treatment with surplus-enhancing features. In other words, while buyers have good reason to be averse to complex products in the SNT markets, they have no reason to be suspicious of or averse to products with many add-ons in the markets with surplus-enhancing add-ons. We believe that this explains why complexity aversion, and the associated fragility of high complexity levels, shows up only in markets with surplus-neutral add-ons but not in markets with surplus-enhancing product features.

Figure 9: The relationship between products' aggregate net feature values and the number of features in the OO Markets with surplus-enhancing and with surplus-neutral features.

(a) Surplus-enhancing add-ons **(b) Surplus-neutral add-ons**

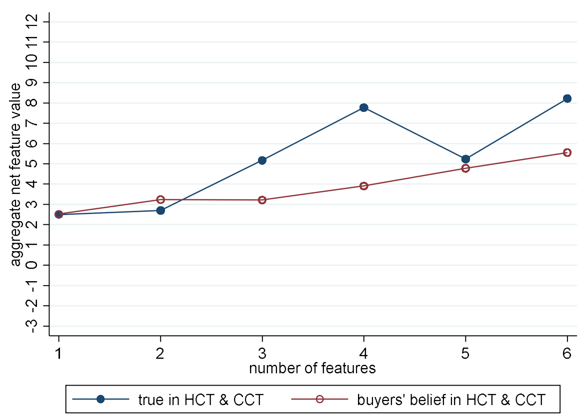

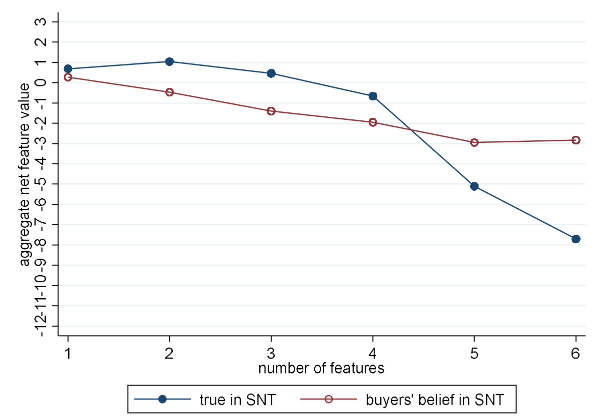

Notes: The figures show the buyers' beliefs about the relationship between products' aggregate net feature values and the number of features and the actual relationship. Figure (a) illustrates the actual and believed relationship in the treatments with surplus-enhancing add-ons (HCT and CCT) while Figure (b) shows the corresponding believed and actual relationships in the treatment with surplus-neutral add-ons (SNT).

5 Summary and Conclusions

While traditional economic models of competitive markets assumed that consumers are able to understand and compare all the products in the market, consumers' attention and information is costly

---

[34] This correlation is also predicted by, e.g., Carlin (2009), Ellison and Wolitzky (2012), and also Chioveanu and Zhou (2013).

and their cognitive capacities are limited. These constraints appear particularly important in modern economies characterized by a flood of products with a large number of different add-on features.

To understand the consequences of obfuscation via add-ons on market prices, competition, and consumer welfare, we designed an experiment with obfuscation opportunities for the sellers and search opportunities for buyers. To identify the causal impact of obfuscation opportunities, we contrast the market with obfuscation opportunities with an otherwise identical control market without obfuscation opportunities. We find that in markets with surplus-enhancing obfuscation opportunities high product complexity becomes a stable market characteristic that severely constrains buyers' product information. We also document that buyers for whom the time needed to assess the overall value of a product is a more convex function of the number of add-ons and those with higher mental cost of search, are likely to avoid studying products deeply and prefer to examine products in a superficial way. Limited product search gives sellers the opportunity to enforce add-on prices that are considerably above marginal cost which allows them to appropriate a substantial share of the total market surplus even in the long run, i.e., even during the final periods of the market.

If we remove obfuscation opportunities price levels quickly converge to marginal cost. In addition, we also find that obfuscation is considerably more fragile when add-ons are surplus-neutral, i.e., when the decision to increase product complexity is independent of the product's overall value. We find, in particular, that buyers are reluctant to buy more complex products in some of the markets with surplus-neutral add-ons, and this reluctance prevails even if we control for all other aspects of the good, thus reducing individual sellers' incentive to provide add-ons. In these markets, product complexity therefore decreases over time and the market approaches competitive conditions. A plausible reason for buyers' reluctance to buy complex products is that higher product complexity is associated with less valuable goods for the buyers in these markets – a fact the buyers realize. On the other hand, complexity aversion is absent when add-on features are surplus enhancing. A potential reason for this is that when add-ons generate, on average, additional surplus, sellers can earn profits while not extracting all the surplus from the add-ons. In this case, products with more add-ons are more valuable for the buyers – which is indeed the case in the OO Markets with surplus-enhancing add-ons. We show that buyers perceive this fact and, therefore, they have little reason to resist buying more complex products.

We believe that these findings may provide a deeper understanding of the forces that sustain obfuscation in markets and may guide theory construction. Documenting the existence of "browsers" and "studiers" and the relationship of these different search styles with individuals' time and mental cost of search may, for example, stimulate new assumptions and theoretical micro-foundations of buyers' behavior in models of obfuscation. In addition, our experimental design may also be useful as a workhorse for studying other exciting questions. Our design could, e.g., be easily extended to study the question whether, and under which conditions, reputation formation among sellers can solve the problem of overpriced and hidden add-ons. Intuitively, reputation formation may solve the problem because sellers who provide reasonably priced high-quality add-ons could have a competitive advantage. At the same time, however, reputation may introduce other competition mitigating forces because it bestows higher price-setting power on high-reputation sellers. Another interesting problem concerns the question, how headline and add-on prices as well as traded quantities in obfuscated markets respond to supply and demand shocks. The sources of price stickiness are of interest for many subfields in economics, and in the case of obfuscated markets the sluggish response of add-on prices to competitive pressures may also translate into sluggish responses to supply or demand shocks. Finally, our design may also be useful as a work horse for studying how different regulatory interventions affect the functioning of obfuscated markets and how different market characteristics that shape the necessity of "add-ons" affect the functioning of markets.[35]

References

Bertini, M., & Wathieu, L. (2008). Research note—Attention arousal through price partitioning. *Marketing Science*, *27*(2), 236-246.

Bourne, R., & Bagley, S. (2023). Junk Fees or Junk Economics? Unpublished working paper, available at: https://www.cato.org/briefing-paper/junk-fees-or-junk-economics

Brown, J., Hossain, T., & Morgan, J. (2010). Shrouded attributes and information suppression: Evidence from the field. *The Quarterly Journal of Economics*, *125*(2), 859-876.

Carlin, B. I. (2009). Strategic price complexity in retail financial markets. *Journal of financial Economics*, *91*(3), 278-287.

Célérier, C., & Vallée, B. (2017). Catering to investors through security design: Headline rate and complexity. *The Quarterly Journal of Economics*, *132*(3), 1469-1508.

Chetty, R., Looney, A., & Kroft, K. (2009). Salience and taxation: Theory and evidence. *American economic review*, *99*(4), 1145-77.

Chiles, B. (2017). Shrouded prices and firm reputation: evidence from the US hotel industry. *Available at SSRN 2952950*.

Chioveanu, I., & Zhou, J. (2013). Price competition with consumer confusion. *Management Science*, *59*(11), 2450-2469.

Choi, M., Dai, A. Y., & Kim, K. (2018). Consumer search and price competition. *Econometrica*, *86*(4), 1257-1281.

Crosetto, P., & Gaudeul, A. (2011). Do consumers prefer offers that are easy to compare? An experimental investigation. *Jena Economic Research Papers*, *44*.

Crosetto, P., & Gaudeul, A. (2017). Choosing not to compete: Can firms maintain high prices by confusing consumers?. *Journal of Economics & Management Strategy*, *26*(4), 897-922.

Diamond, P. A. (1971). A model of price adjustment. *Journal of economic theory*, *3*(2), 156-168.

Dufwenberg, M., & Gneezy, U. (2000). Price competition and market concentration: an experimental study. *international Journal of industrial Organization*, *18*(1), 7-22.

Ellison, G. (2005). A model of add-on pricing. *The Quarterly Journal of Economics*, *120*(2), 585-637.

Ellison, G., & Ellison, S. F. (2009). Search, obfuscation, and price elasticities on the internet. *Econometrica*, *77*(2), 427-452.

Ellison, G., & Ellison, S. F. (2018). Search and obfuscation in a technologically changing retail environment: Some thoughts on implications and policy. *Innovation Policy and the Economy*, *18*(1), 1-25.

Ellison, G., & Wolitzky, A. (2012). A search cost model of obfuscation. *The RAND Journal of*

[35] We owe this suggestion to one reviewer who pointed out that for some markets (such as financial markets) add-ons appear unavoidable in the sense that for any investment, one needs to check counterparty risks or additional fees beyond the ones that are regulated. It would thus be interesting to study such a "search-intensive" market in which sellers must choose *all* aspects of a product, and the consumers need to click on each one to see whether the "add-on" is offered and at what price.

Economics, *43*(3), 417-441.

Fischbacher, U. (2007). z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental economics*, *10*(2), 171-178.

Gabaix, X., & Laibson, D. (2006). Shrouded attributes, consumer myopia, and information suppression in competitive markets. *The Quarterly Journal of Economics*, *121*(2), 505-540.

Gaudeul, A., & Sugden, R. (2012). Spurious complexity and common standards in markets for consumer goods. *Economica*, *79*(314), 209-225.

Greenleaf, E. A., Johnson, E. J., Morwitz, V. G., & Shalev, E. (2016). The price does not include additional taxes, fees, and surcharges: A review of research on partitioned pricing. *Journal of Consumer Psychology*, *26*(1), 105-124.

Grubb, M. D. (2015). Failing to choose the best price: Theory, evidence, and policy. *Review of Industrial Organization*, *47*(3), 303-340.

Gu, Y., & Wenzel, T. (2015). Putting on a tight leash and levelling playing field: An experiment in strategic obfuscation and consumer protection. *International Journal of Industrial Organization*, *42*, 120-128.

Heidhues, P., Kőszegi, B., & Murooka, T. (2012). Deception and consumer protection in competitive markets. *The Pros and cons of consumer Protection*, *44*.

Heidhues, P., Kőszegi, B., & Murooka, T. (2016). Inferior products and profitable deception. *The Review of Economic Studies*, *84*(1), 323-356.

Heidhues, P., & Kőszegi, B. (2017). Naivete-based discrimination. *The Quarterly Journal of Economics*, *132*(2), 1019-1054.

Heidhues, P., Johnen, J., & Kőszegi, B. (2021). Browsing versus studying: A pro-market case for regulation. *The Review of Economic Studies*, *88*(2), 708-729.

Hefti, A., Liu, S., & Schmutzler, A. (2020). Preferences, confusion and competition. *University of Zurich, Department of Economics, Working Paper*, (344).

Hossain, T., & Morgan, J. (2006). Shrouded Attributes and Information Suppression: Evidence from Field Experiments.

Huck, S., Normann, H. T., & Oechssler, J. (2004). Two are few and four are many: number effects in experimental oligopolies. *Journal of economic behavior & organization*, *53*(4), 435-446.

Jin, G. Z., Luca, M., & Martin, D. J. (2018). *Complex disclosure* (No. w24675). National Bureau of Economic Research.

Kalaycı, K. (2015). Price complexity and buyer confusion in markets. *Journal of Economic Behavior & Organization*, *111*, 154-168.

Kalaycı, K. (2016). Confusopoly: competition and obfuscation in markets. *Experimental Economics*, *19*(2), 299-316.

Kalaycı, K., & Potters, J. (2011). Buyer confusion and market prices. *International Journal of Industrial Organization*, *29*(1), 14-22.

Kalaycı, K., & Serra-Garcia, M. (2016). Complexity and biases. *Experimental Economics*, *19*(1), 31-50.

Kim, H. M. (2006). The effect of salience on mental accounting: how integration versus segregation of payment influences purchase decisions. *Journal of behavioral decision making*, *19*(4), 381-391.

Miravete, E. J. (2013). Competition and the use of foggy pricing. *American Economic Journal: Microeconomics*, *5*(1), 194-216.

Morwitz, V. G., Greenleaf, E. A., & Johnson, E. J. (1998). Divide and prosper: consumers' reactions to partitioned prices. *Journal of marketing research*, *35*(4), 453-463.

Muir, D., Seim, K., & Vitorino, M. A. (2013). *Price obfuscation and consumer search: An empirical analysis* (pp. 1-48). Working paper.

Piccione, M., & Spiegler, R. (2012). Price competition under limited comparability. *The quarterly journal of economics*, *127*(1), 97-135.

Richards, T. J., Klein, G. J., Bonnet, C., & Bouamra-Mechemache, Z. (2019). Strategic obfuscation and retail pricing. *Review of Industrial Organization*, 1-31.

Santana, S., Dallas, S. K., & Morwitz, V. G. (2020). Consumer reactions to drip pricing. *Marketing Science*, *39*(1), 188-210.

Shulman, J. D., & Geng, X. (2013). Add-on pricing by asymmetric firms. *Management Science*, *59*(4), 899-917.

Seim, K., Vitorino, M. A., & Muir, D. M. (2017). Do consumers value price transparency? *Quantitative Marketing and Economics, 15*(4), 305-339.

Sims, C. A. (2003). Implications of rational inattention. *Journal of monetary Economics*, *50*(3), 665-690.

Spiegler, R. (2006). Competition over agents with boundedly rational expectations. *Theoretical Economics*, *1*(2), 207-231.

Spiegler, R. (2014). Competitive framing. *American Economic Journal: Microeconomics*, *6*(3), 35-58.

Sugden, R., & Zheng, J. (2018). Do consumers take advantage of common pricing standards? An experimental investigation. *Management Science*, *64*(5), 2126-2143.

Völckner, F., Rühle, A., & Spann, M. (2012). To divide or not to divide? The impact of partitioned pricing on the informational and sacrifice effects of price. *Marketing Letters*, *23*(3), 719-730.

Weitzman, M. L. (1979). Optimal search for the best alternative. *Econometrica: Journal of the Econometric Society*, 641-654.

Wilson, C. M. (2010). Ordered search and equilibrium obfuscation. *International Journal of Industrial Organization*, *28*(5), 496-506.

Wolinsky, A. (1983). Prices as signals of product quality. *The review of economic studies*, *50*(4), 647-658.

Xia, L., & Monroe, K. B. (2004). Price partitioning on the internet. *Journal of Interactive marketing*, *18*(4), 63-73.

Online Appendices

Table of Content

Appendix 1
Examples of Screens shown to Buyers and Sellers

Figure A1: example screens in the buyer stage in market with obfuscation opportunity

(a) Buyer screen before buyer searched though add-ons

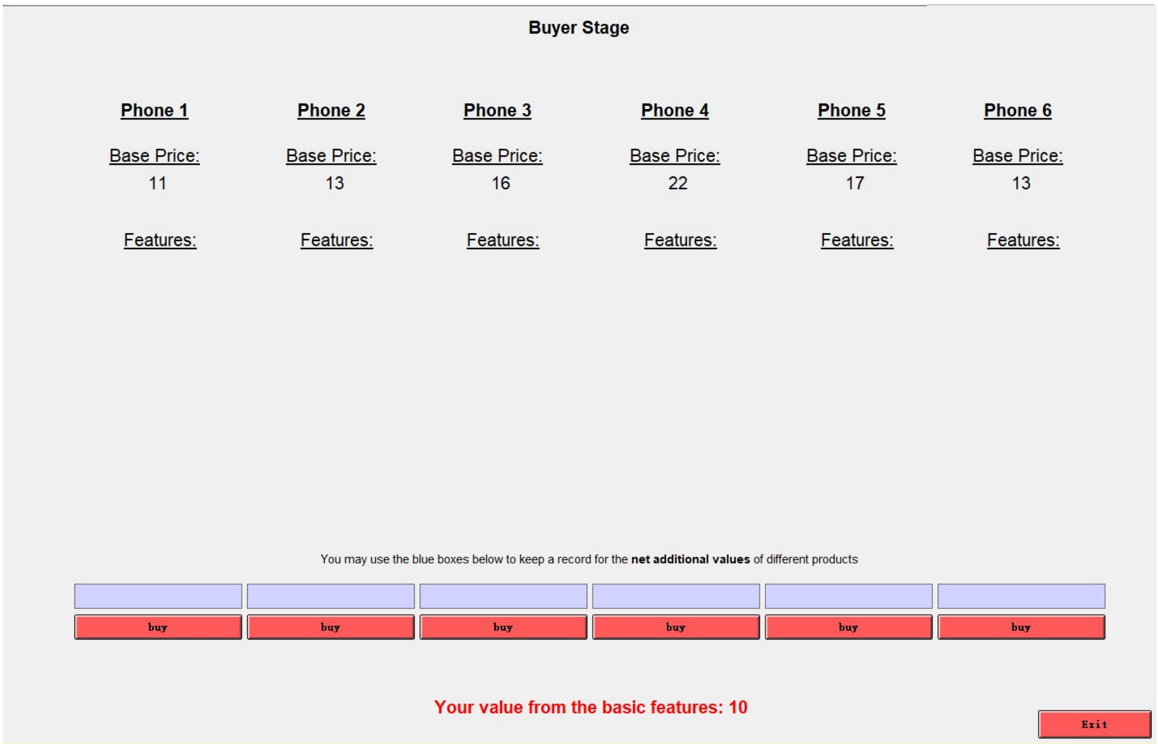

(b) Buyer Screen after searching the first add-on of Phone 1

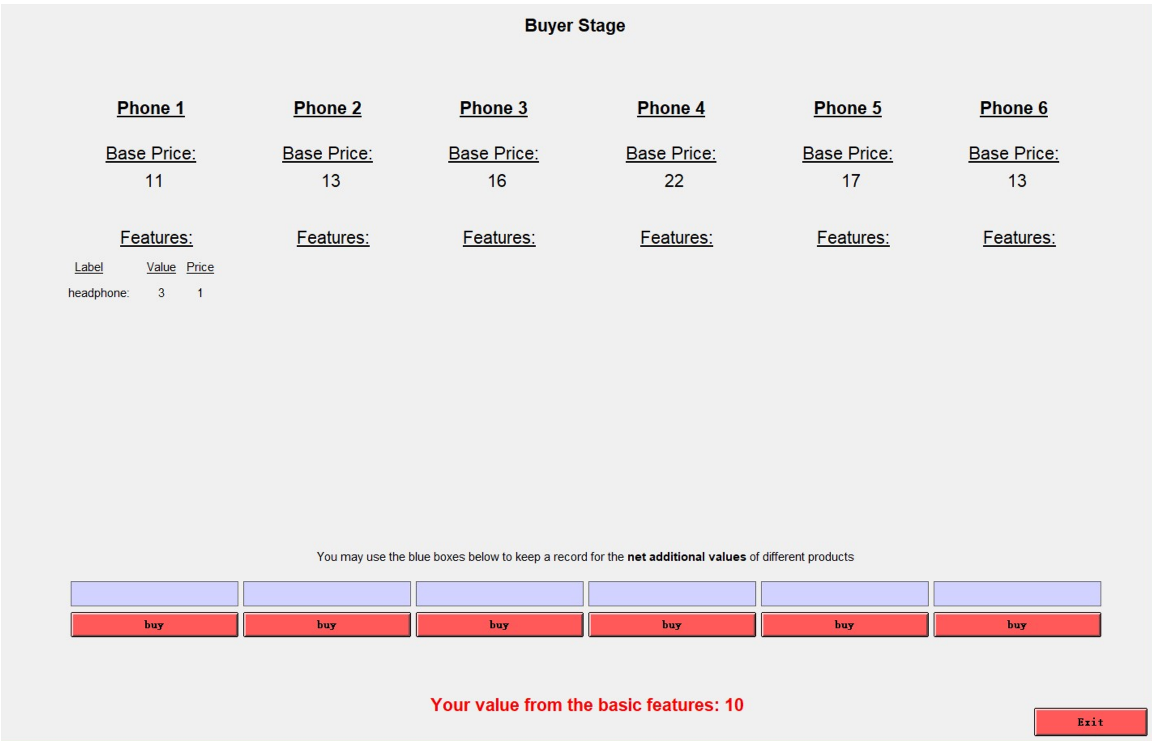

(c) Buyer Screen after searching the second add-on of Phone 1

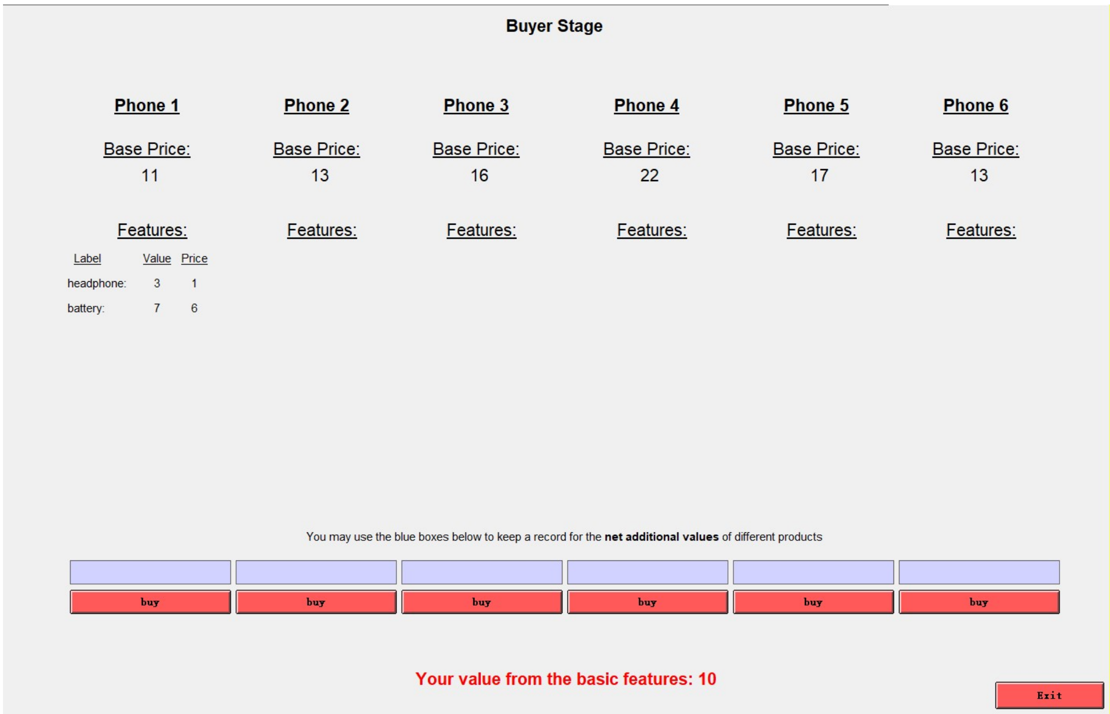

Notes: In the buyer stage, buyers first immediately only see (a). With each click on, e.g., Phone 1, buyers can see one more feature of that product like in (b) and (c). If buyers click on another product, then the features of this other product are also shown one by one as buyers click.

Figure A2: example screens of the feedback stage

(a) Feedback screen for seller of phone 5

(b) Feedback screen for a buyer

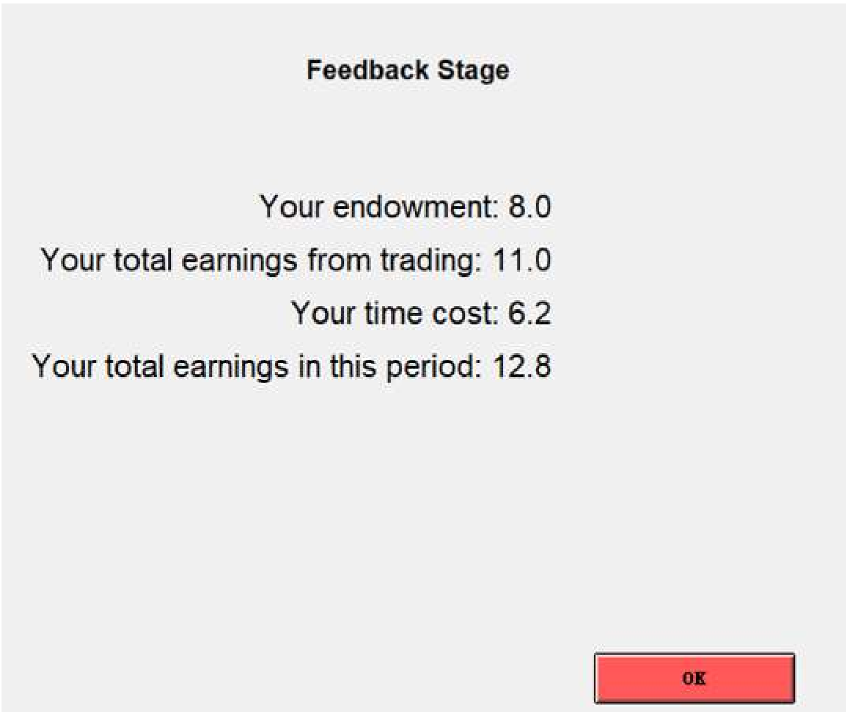

Figure A3: an example screen at the buyer stage in the NO Market of the main treatments described in Table 1

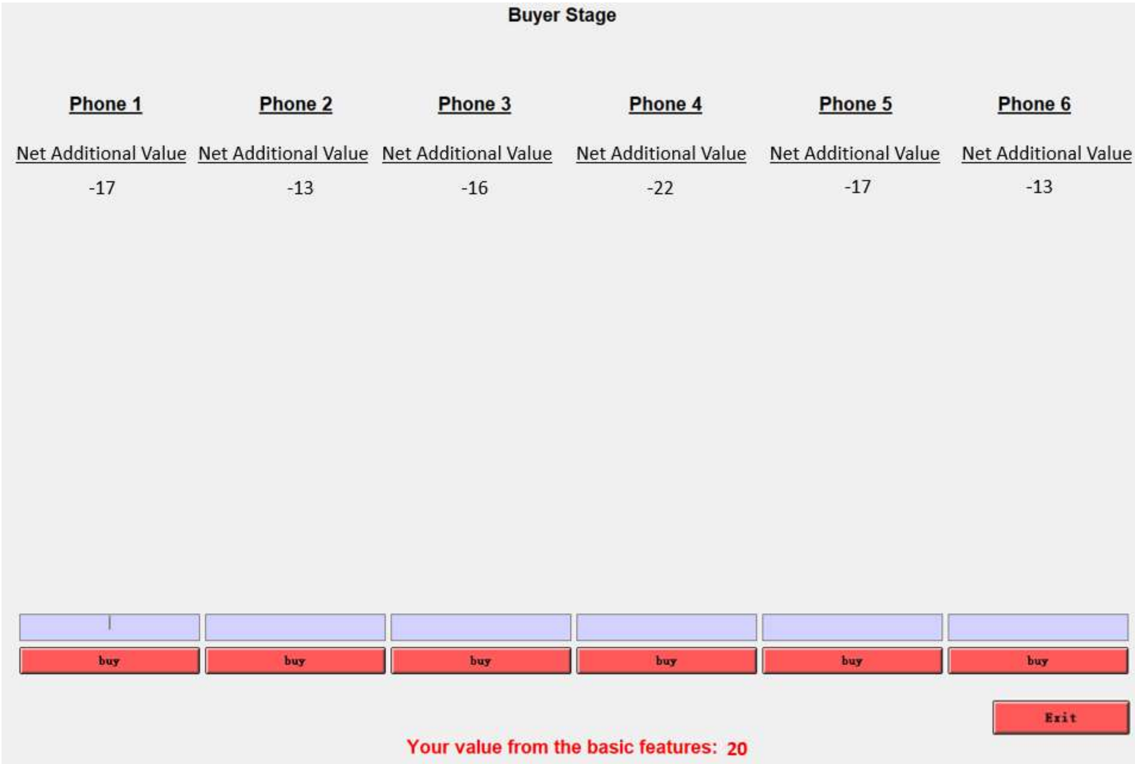

Notes: The overall net value v_o of products is described with the name "net additional value", so "net additional value" is shown on the buyer's stage in the NO Market.

Figure A4: an example screen at the buyer stage in the NO Market of the robustness check treatments described in Table 1

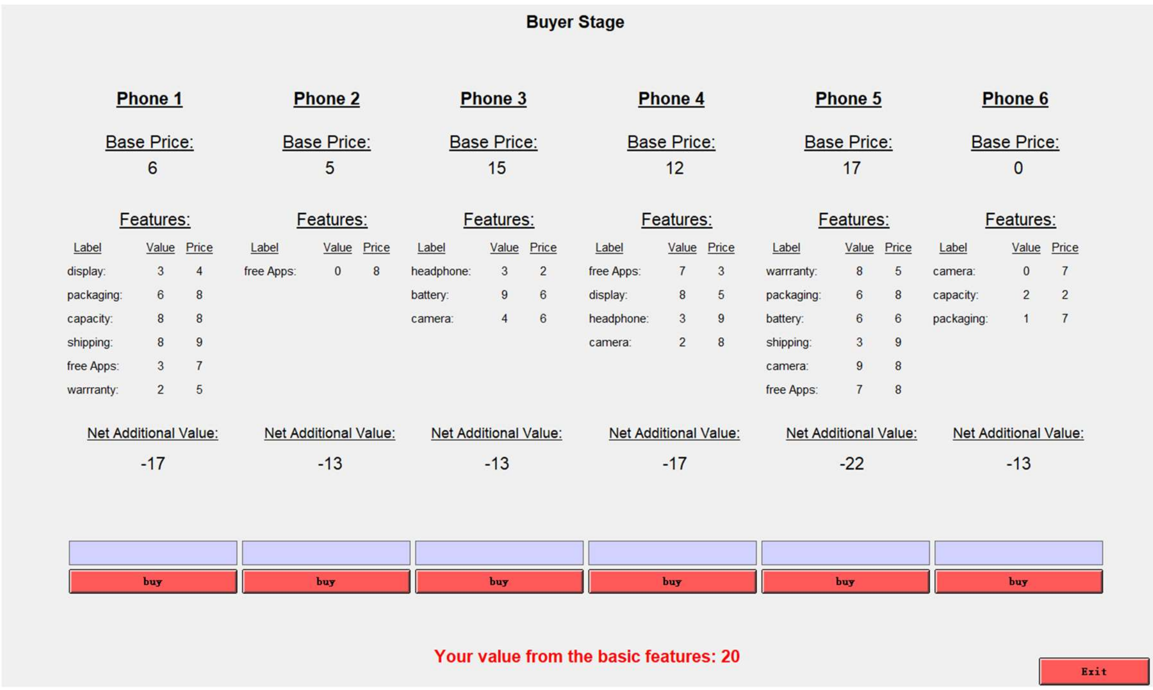

Figure A5: aggregate feature values, feature prices and feature costs of offered products in the OO markets

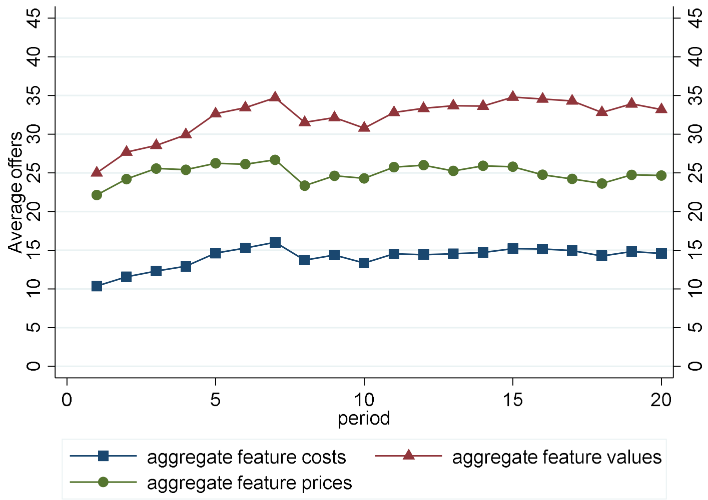

Appendix 2
Experimental tasks to measure the time needed and the mental cost of computing the overall net value of products

In this appendix, we explain the two tasks each buyer completed in Part 3 of the experiment. These tasks were designed to measure two key cost components associated with product search: (i) the *time needed* for correctly assessing (i.e., computing) the overall net value v_o of products with varying levels of complexity, and (ii) the *mental cost* of such computations, which reflects how much buyers value avoiding the mental effort required for such computations.

In **Task 1**, we measured each buyer's time needed for computing the overall net value v_o of products with varying complexity. v_o is given by the product's aggregate feature value minus its base price and aggregate feature price: $v_o \equiv v_f - p_b - p_f$. Each buyer was given an endowment of 40 ECU and was asked to compute the overall net value of 14 products. During the task, buyers faced the same time cost as during the market interactions (i.e., 0.1 ECU per second) ensuring that the task closely resembled the market environment with similar incentives to economize on time. Like in the market environment, buyers also had an incentive to compute the overall net value correctly because if a buyer's computation was incorrect, he or she was required to try again until the correct value was obtained. Buyers' final payment for Task 1 was determined by subtracting their total monetary time cost from their 40 ECU endowment.

The 14 products in Task 1 varied in the number of add-on features. There were two products with 0 features, two products with 1 feature, two products with 2 features, and so on, up to two products with 6 features. All buyers computed the same set of 14 products, but the order in which the products were presented was randomly determined for each buyer. This setup allowed us to measure buyers' time needed *per add-on* for computing v_o for products with different levels of complexity. We were particularly interested in examining whether the time needed *per add-on* is rising in the number of features.

In **Task 2**, we measured buyers' mental cost of computing overall net values. Buyers were asked to compute the overall net values of another set of 14 products, but they were also given the option to let the computer perform all computations for them. To measure each buyer's mental cost of computing v_o for the 14 products, we elicited the highest amount of ECUs they were willing to pay (WTP) for delegating the computations to the computer. The BDM (Becker, DeGroot and Marschak, 1964) mechanism was used to ensure that this WTP measure was incentive compatible.[36] The fact that Task 2 was conducted after Task 1 ensures that buyers had directly experienced the computation task before and thus understood how much mental efforts was involved in these computations.

One concern may be that buyers with higher computation ability might report lower WTP simply because they need less time for computing the values. To address this concern, we designed the payment structure in Task 2 as follows. If the realization of the random BDM price implied that the buyer had to compute the values themselves, their earnings in Task 2 were the 40 ECU endowment minus their total

[36] That is, a random price for this option was drawn. If the WTP reported by a subject was higher than the random price, then the subject paid the price and the computer correctly computed v_o for all 14 products. But if the WTP of a subject was lower than the random price, then the subject did not pay the price and had to compute v_o for all 14 products by himself/herself.

monetary time cost in Task 2. If, instead, the realization of the random BDM price implied that the computer handled the computations, their earnings were the 40 ECU endowment minus their total time cost from Task 1 and the random price chosen by the BDM mechanism. As the total time cost from Task 1 is the best available proxy for each buyer's computation ability, computationally more skilled buyers also earn more from delegating the computations to the computer, i.e., the announced WTP measure controls for buyers' (heterogeneous) computational abilities.

Appendix 3
The role of add-on labels and the number of add-ons for buyers' and sellers' behavior

In this appendix, we ask the following three questions: (i) Is buyers' behavior affected by add-on labels? (ii) Do sellers believe that buyers have a non-pecuniary preference for specific add-on labels? (iii) Do sellers believe that buyers prefer per se a higher number of add-ons?

As described in the experimental design section, each add-on feature of a product involves 3 parts: a feature label, a feature value, and a feature price. While feature values and feature prices are obviously relevant for both buyers' and sellers' payoffs and decisions, it is unclear whether the feature labels also play a behavioral role. One possibility is that buyers and sellers largely neglect the labels and concentrate merely on the pecuniary aspects (e.g., the values and prices) of the goods. The other possibility is that buyers exhibit a preference for certain labels over others, despite the fact that labels did not affect their pecuniary payoffs; if sellers anticipate such a preference, they may have an incentive to increase the likelihood that certain add-on labels are mentioned in the list of add-ons.

In this appendix we conducted several analyses that examine the role of labels for buyers' and sellers' behavior. Based on these analyses below, we conclude that the labels of add-on features had little or no impact on buyers' and sellers' behavior. Thus, as we intended with our experimental design, the labels just provided a more realistic context but had no behavioral influence.

Buyer behavior with regard to add-on labels

How did the buyers themselves view the role of feature labels? To answer this question, we explicitly elicited buyers' views in a survey conducted in several experimental sessions where they could indicate their (dis)agreement on a 5-point Likert scale (1 = strong disagreement; 2 = disagreement; 3 = neutral; 4 = agreement; 5 = strong agreement) with the statement "T*he labels were important for my purchase decisions*". For this question the buyers expressed almost unanimously a "strong disagreement" with an average value of 1.15. In contrast, the buyers expressed almost universal agreement (value of 4.90) for the question: "*I paid little attention to the labels of the extra features because they were irrelevant for my monetary payoff*". Finally, we also asked buyers the question "*I only cared for the monetary net value (= feature value – feature price) of the extra features because they affected my earnings.*" Here, the buyers also expressed a strong agreement with a score of 4.81. Thus, buyers surevey answers suggest that the they basically cared only about their monetary payoff while the labels did not affect their behavior.

In total, there are 10 possible labels that may appear for an add-on feature in the experiment, and, for the sake of the analyses conducted below, we code them as follows: 1 = "shipping"; 2 = "display"; 3 = “capacity"; 4 = "camera"; 5 = "weight"; 6 = "headphone"; 7 = "battery"; 8 = "packaging"; 9 = "warranty"; 10 = "free Apps”.

We also performed three sets of analyses based on data from the surplus-enhancing main treatments (HCT and CCT). First, we ran 10 t-tests (clustered at the market level) on whether buyers’ purchase decisions differed depending on whether a product contains a particular label (e.g. Label 1) or not. None of these 10 t-tests is significant at the 5% level (p = 0.430, p = 0.706, p = 0.565, p = 0.245, p = 0.210, p = 0.264, p = 0.704, p = 0.732, p = 0.171, and p = 0.097, respectively).

Second, we added a dummy variable for each of the 10 labels that took a value of 1 if the label was present in a product and took a value of zero otherwise. We included these dummies into the regressions of buyers’ purchase decisions reported in Table 2 of the paper, i.e., we examine the impact of each label while controlling for all the other regressors mentioned in Table 2. There are two different regression specifications in Table 2; but for both specifications, none of the 10 labels shows a significant impact on buyers’ purchase decisions (p = 0.488, p = 0.778, p = 0.425, p = 0.093, p = 0.179, p = 0.537, p = 0.532, p = 0.746, p = 0.644, and p = 0.175 for the first specification; p = 0.525, p = 0.885, p = 0.442, p = 0.117, p = 0.237, p = 0.725, p = 0.365, p = 0.697, p = 0.771, and p = 0.331 for the second specification).

In a **third** type of analysis, we examined whether certain combinations of labels played a role in buyers’ purchase behavior. For this purpose, we ran 45 regressions (one for each possible two-label combination of the 10 labels) with buyers’ demand for products as the dependent variable, a dummy for label X, another dummy for label Y and an interaction term (XY) capturing the joint influence of the respective labels. Among these 45 regressions, the individual dummies were always insignificant, and in only 2 regressions we observed a significant coefficient for the interaction term at the 5% level, which is what one would expect to occur by chance because $45 \times 0.05 = 2.25$. Thus, we also find no evidence that two-label combinations played any role in buyers’ purchase behavior.

Finally, the behavioral evidence from the NO Markets also strongly suggests that the labels were irrelevant for buyers purchase behaviors. Recall that in the NO Markets of the main treatments the buyers could only see the overall net value (v_o) of all offered products. We implemented the NO Markets of the main treatments in this way because we assumed that buyers were only interested in the pecuniary payoff and not in other aspects such as specific labels. We find that in only 3% of the cases the buyers in the NO Markets of the main treatment with surplus-enhancing add-ons do *not* maximize their overall net value. In the NO Markets of the robustness check treatments the buyers do not only see the overall net value of each offered product but they also see all the add-on features (with labels, values, and prices) of each available product in the market. Thus, if buyers indeed have an independent preference for specific labels or for the number of add-ons per se, we should see a higher number of buyer decisions that does not maximize their overall net value compared to the NO Markets of the main treatment. In other words, buyers should be willing to sacrifice pecuniary payoff for the sake of getting products with specific labels or products with a higher number of add-ons. It turns out, however, that there is only a negligible rise (by 1.5 percentage points) in the percentage of cases in which buyers do not maximize their overall net value. Thus, buyer behavior in NO Markets also suggests that buyers were primarily interested in their own pecuniary payoff while labels did not matter.

Taken together, neither the survey evidence nor the econometric evidence on buyers' purchase behavior in the OO Markets nor buyer behavior in the NO Markets provides an indication that the labels mattered.

Sellers' beliefs and behaviors with regard to add-on labels

Although buyers may not exhibit a preference over the labels, sellers may believe buyers have such preferences and design their products accordingly. To address this issue, we elicited sellers beliefs about the role of labels for buyers' purchase behavior in the surveys conducted in several experimental sessions. In particular, we elicited sellers' (dis)agreement with the following two statements (1 = strong disagreement; 2 = disagreement; 3 = neutral; 4 = agreement; 5 = strong agreement): "*I believe that the add-on labels were important for buyers' purchase decisions.*" And: "*I believe that buyers only cared for the monetary net values (= feature value – feature price) of the extra features because that affected their earnings.*" The sellers indicated that they disagreed with the first statement (value of 1.92), indicating that they thought that the labels were not important for the buyers' decisions. The same conclusion follows also from the strong agreement (value of 4.48) with the second statement. We also had a third statement on this issue that implies the same conclusion: "*I believe that buyers paid little attention to the labels of the extra features because they were irrelevant for their monetary payoff.*" Sellers strongly agreed with this statement (value of 4.50).

We also performed three analyses to investigate whether potential beliefs about the feature labels' impact on buyers' behavior affected sellers' behavior. First, if sellers believe that certain feature labels enhance the probability that buyers will buy their product, they have an incentive to re-randomize the assignment of feature labels, similar to what they did for the net feature values of their add-ons (see Figure 6 in main text). Why? Because, as shown in Section 3.3, the sellers clearly anticipated that many buyers will not examine the add-on features that are further down the list. Thus, the more preferred feature labels should appear more frequently at the initial add-on positions. However, if we compute the feature labels that appear on average across the 6 different add-on positions using the coding of the 10 labels mentioned above (1 = "shipping"; 2 = "display"; 3 = "capacity"; 4 = "camera"; 5 = "weight"; 6 = "headphone"; 7 = "battery"; 8 = "packaging"; 9 = "warranty"; 10 = "free Apps"), then we find that the average label codes are 5.53, 5.45, 5.44, 5.61, 5.35, 5.37 across the 6 add-on positions. These average label codes are all roughly what one would expect if labels have been assigned by chance, in which case the average should be 5.5. Thus, deliberate manipulation of labels through re-randomization is absent.

We also performed a second analysis, where we simply compute the relative frequency with which each label appears in any feature position of the products. That is, we simply take the *universe of all the single add-ons* (recall that each add-on contains a feature label, feature value and feature price) that appeared in any product offered during the 20 periods and count the relative frequency with which each of the 10 labels appeared. Thus, if none of the 10 labels is favored in sellers' decisions, they should all appear with similar frequencies, which is what we find: the 10 labels appear with frequencies of 9.24%, 10.91%, 10.08%, 10.38%, 10.15%, 11.29%, 9.92%, 9.47%, and 9.70%, respectively.

In a third analysis we take the *universe of all the products* offered during the 20 periods and count the relative share of products for which a given label appeared. If none of the 10 labels is favored in sellers' decisions, the relative share of products that contains a given label should be identical across the ten labels, which is again very close to what we find: the 10 labels appear with frequencies of 50.01%, 50.53%, 53.41%, 52.95%, 50.76%, 51.67%, 52.27%, 46.21%, 51.21%, and 50.30%, respectively. In view of the fact that the average product contains roughly 5 features, each label has a roughly 50% chance of being included in a product.

Buyer behavior with regard to the number of add-ons

As already shown and discussed in Tables 2 and 3 of the main text, we find that buyers do not exhibit a significant non-pecuniary preference in favor of products with more add-ons, i.e., if we control for a product's (and it's best competitor's) values and prices, the number of add-ons does not increase product demand. Moreover, in the low complexity markets of the Surplus Neutral Treatment we even find that buyers exhibit a significant aversion to products with more add-on features, which may be an important counterforce for sellers' decisions to make their products complicated.

Finally, the evidence from the NO Markets also suggests that the number of add-ons are, per se, irrelevant for buyers' purchase decisions. Recall that in the main treatments of the NO Markets the buyers could only see the net additional value of each product while in the robustness check treatments they could see the add-ons of all products. Thus, if buyers indeed cared for the number of add-ons per se, one would expect more deviations from the money maximizing purchase decision in the robustness check treatments. However, we observed only a negligible increase in the share of non-money maximizing buyer decisions in the robustness check treatment of the NO Markets. Thus, as also confirmed by buyers' answers to the surveys mentioned above, they almost exclusively cared for the monetary payoff they receive from the purchased product while the add-on labels and the number of add-ons were irrelevant for their purchase decisions.

Seller beliefs with regard to buyers' preferences about the number of add-ons

The survey in Part 3 asked the sellers about their beliefs regarding the buyers' preferences between a product A with 4 add-ons and a product B with 6 add-ons but product A gives buyers a higher monetary gain by two units. Thus, if sellers believe that buyers prefer product B, they believe that buyers have a non-pecuniary preference for more add-ons. However, 92% of sellers believe that buyers prefer A, and only 4% believe that buyers are indifferent or prefer B (4%). Thus, these answers again indicate that sellers assume that buyers basically care only for the money value of the offered products.

Appendix 4
The relationship between base price and aggregate net feature values

In this appendix, we examine whether lower base prices are associated with worse add-on features in the sense that the net aggregate value of add-on features is positively associated with base prices. Such a relationship could trigger buyers' suspicion against very low base prices. The regression below and the associated Figure A 5 indeed indicate a positive association.

Table A1: Regression of net aggregate feature values on base prices in OO markets

	Net aggregate feature values
Base price	0.43** (0.16)
Period dummies for period 2-20	√
Constant	√
No. of observations	1320
R-square	0.11

Notes: Standard errors are clustered on the market level. *** indicates significance at the 0.01 level, ** at the 0.05 level, and * at the 0.10 level.

Figure A5: The relation between products' aggregate net feature values and their base prices

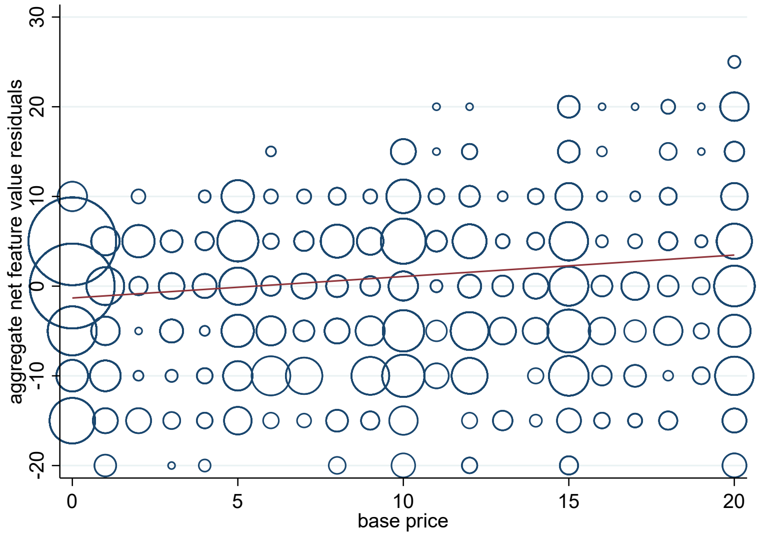

Notes. The figure shows the scatter plot of the residual aggregate net feature values as a function of base prices after controlling for period dummies. The fitted line represents the regression coefficient, and the bubbles represent the smoothed frequency at each data point. The figure is based on the pooled data from the Half-Cost Treatment and the Convex Cost Treatment.

To what extent do buyers indeed shy away from the lowest base prices? To answer this question, we examined the influence of base prices on buyers' first visited products and on the traded products in Figure A6a and A6b below. Figure A6a depicts the average base prices of buyers' *first visited products*, along with the minimum and the maximum base prices during a period. We see that the buyers first visit products with relatively low base prices, but the figure also shows that there is a significant difference ($p < 0.001$) between the lowest base price and the average base price buyers first visit.

In view of the fact that many buyers (37%) visit only one product, the first visited product is likely to have an impact on the actually traded products. In fact, Figure A6b illustrates that the average base price *in traded products* is also higher than the lowest base price in the market. Note also that the patterns displayed in Figures A6a and A6b prevails not just "on average" but are prevalent throughout the whole 20 periods. Thus, a nonnegligible number of buyers shy away from buying the product with the lowest base price which suggests some constraints on competition via base prices.

Figure A6: Base prices and buyers' search behavior in the OO Market

(a) Average Base price of first visited product (b) Average base price in traded products

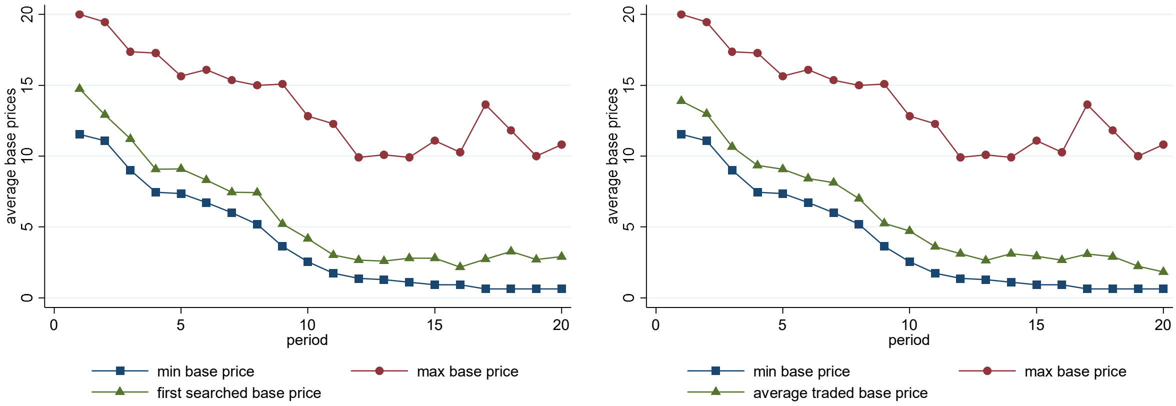

Note*:* The figures show how buyers' search and purchase behavior is affected by sellers' base prices. **(a)** The average base price of buyers' first visited products together with the range of base prices during a period. The average base price of the first visited products is higher than the lowest base prices in the market with p < 0.001. **(b)** The average base price of traded products together with the range of base prices. The average base price of traded products is higher than the lowest base price in the market with p < 0.001 for the entire 20 periods and with p = 0.006 for period 16-20.

Appendix 5

How does an increase in search costs affect behaviors and market outcomes

We discussed the empirical effects of a search cost increase in Section 3.4 without explicitly showing the associated graphs and tables. Below we provide a presentation of these results in terms of graphs and tables:

(i) In each period, the time spent searching in the high-cost market is lower compared to the low-cost market, and the search cost increase reduces the average time spent searching in the market significantly ($p < 0.004$) by 34%. Figure A.7 below illustrates this result graphically.

Figure A.7: Time spent searching in the market

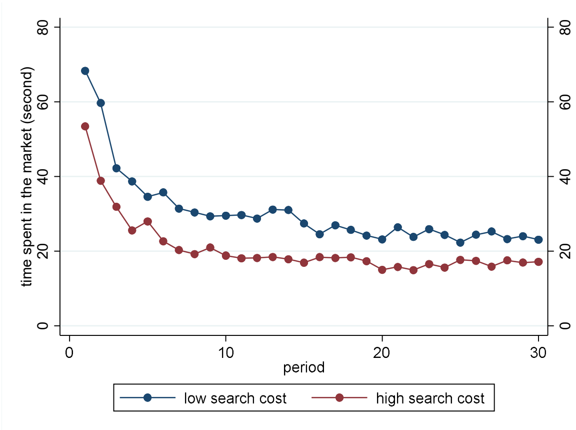

(ii) The search cost increase shifts the whole distribution of individuals' depth of search to the left and, on average, the depth of search is reduced by 20.9% ($p < 0.036$). Likewise, the whole distribution of the breadth of search shifts to the left and the average breadth is reduced by 13.5% ($p < 0.020$). Figure A.8 and A.9 below illustrate these results graphically.

Figure A8: The distribution of individuals' depth of search

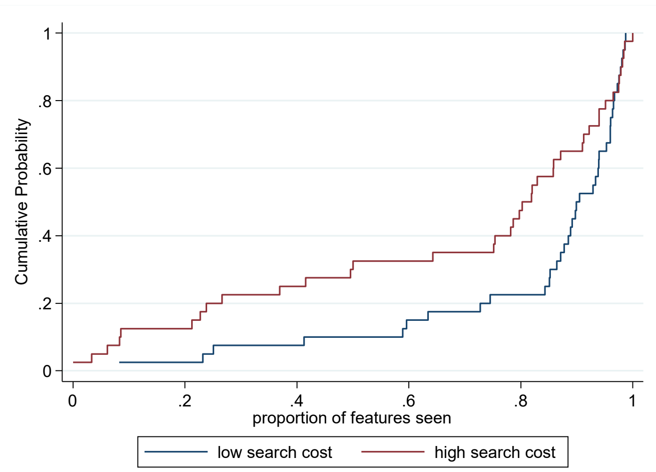

Figure A9: The distribution of individuals' breadth of search

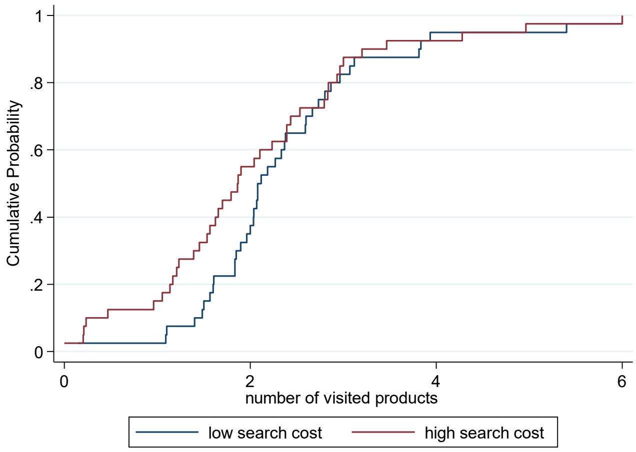

(iii) Higher search costs generate a large reduction in buyers' responsiveness to add-on prices. In particular, at higher search cost the buyers' response to a rise in a product's aggregate feature price (aggregate feature value) is reduced by 59% (56%). Likewise, the buyers' response to the aggregate feature price (aggregate feature value) of the product's best competitor is even reduced by 71% (64%). Table A2 below shows these results in detail:

Table A2 below shows that a rise in a product's own aggregate feature values by 10 units increases sales by 1.6 units in the low search cost treatment while a rise in a product's own aggregate feature prices by 10 units reduces sales by 1.7 units. The interaction terms with the high search cost dummy indicate that these responses are reduced by 0.9 (i.e., 56%) for values and by 1.0 (59%) for prices.

A rise in the best competitor's aggregate feature value (price) by 10 units reduces (increases) the sales of a product by 1.4 units in the low-cost treatment according to Table A2. These responses are, however, again strongly mitigated in the high search cost treatment; the response to the best competitor's aggregate feature values is reduced by 0.9 (64%) units and the response to the best competitor's aggregate feature prices is smaller by 1.0 (71%) units.

Table A2: Buyers' responses to a product's own and the competing products' characteristics in low-cost and high-cost treatments

Dependent Variable	Units sold
Base price	-0.24*** (0.01)
Base price × high search cost	0.05 (0.03)
Aggregate feature value	0.16*** (0.01)
Aggregate feature value × high search cost	-0.09*** (0.02)
Aggregate feature price	-0.17*** (0.02)
Aggregate feature price × high search cost	0.10*** (0.02)
Best competitor's base price	0.17*** (0.00)
Best competitor's base price × high search cost	-0.09** (0.03)
Best competitor's aggregate feature value	-0.14*** (0.01)
Best competitor's aggregate feature value × high search cost	0.09*** (0.02)
Best competitor's aggregate feature price	0.14*** (0.01)
Best competitor's aggregate feature price × high search cost	-0.10*** (0.02)
Number of features	-0.11* (0.05)
Number of features × high search cost	-0.02 (0.09)
Average number of features among competitors	0.09 (0.15)
Average number of features among competitors × high search cost	-0.05 (0.14)
Period and Period × high search cost	√
Constant and Constant × high search cost	√
No. of observations	1439
R-square	0.27

(iv) Higher search costs increase the sellers' share of the surplus from add-ons from 35.3% to 58.9% (p = 0.057), indicating that sellers are able to enforce a considerably higher mark-up over aggregate feature cost.

Figure A10 below shows that the aggregate feature costs are roughly centered around 10 ECUs while the aggregate feature values are typically at or slightly above 30 ECUs in both markets. However, while the aggregate feature prices are (from periods 10 onwards) typically around 15 ECUs in the low search cost market, they are between 20 and 25 ECUs in the high-cost market. Thus, the sluggish buyer responses to the characteristics of add-on features enables the sellers to enforce considerably higher aggregate feature prices in the high-cost market, which explains their higher share of the surplus.

Figure A10: Buyer surplus from add-on features

Low search cost

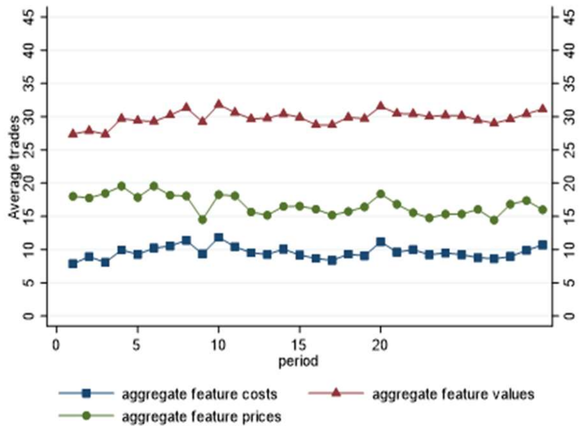

High search cost

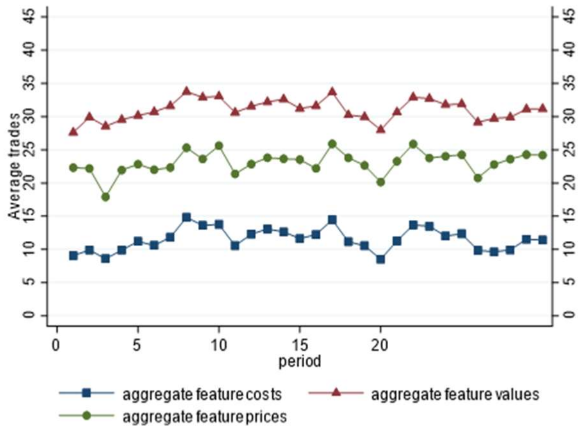

Appendix 6

Robustness with regard to number of periods

To what extent did the market in our OO Markets already converge to stable behavioral outcomes in periods 16-20? To examine this question, we ran the robustness check treatments for 30 periods and compared the behavioral and market outcomes in periods 16-20 with those in periods 26-30. If the period 26-30 outcomes are statistically indistinguishable with those in periods 16-20, we have little reason to assume that markets have not yet converged in periods 16-20.

It turns out that buyers' and sellers' behaviors as well as relevant market outcomes between periods 16-20 and periods 26-30 are very similar and statistically indistinguishable, suggesting that the market reached stable behavioral patterns already in periods 16-20. Moreover, this stability occurs both in the high- and the low-cost treatment.

For example, the time buyers spent searching in the market during a period differs on average only 0.02 seconds (0.46 seconds) in periods 26-30 of the high (low) cost treatment compared to periods 16-20; and these small differences across time periods are clearly insignificant with $p > 0.8$ based on t-tests clustered at the market level. The similarity of time spent per period in the market across the two time phases (16-20 versus 26-30) is also visible in Figure A.7 of the previous appendix. Likewise, buyers' search breadth in terms of products visited in the market is very similar and statistically indistinguishable between periods 16-20 and 26-30 with $p=0.740$ for the high-cost treatment and $p=0.878$ in the low-cost treatment. Finally, the same picture emerges with regard to search depth with $p=0.522$ ($p=0.169$) in the high (low) cost treatment.

Sellers' behavior – in terms of the average number of add-ons per period and the average buyer surplus from add-ons offered – is also very stable across periods 16-20 versus 26-30. Typically, sellers offer on average between 4 and 5 add-ons, and this is not different between the two phases both in the high-cost treatment ($p=0.711$) and the low-cost treatment ($p=0.131$). Likewise, the average offered buyer surplus from add-ons is very stable across time so that no significant differences between periods 16-20 and 26-30 exist ($p=0.142$ in the high-cost treatment; $p = 0.702$ in the low-cost treatment).

Finally, market outcomes – in terms of the *actually traded* buyer surplus *from add-ons* and the *actually traded total* buyer surplus – also do not differ significantly across these periods. For the traded buyer surplus from add-ons the respective p-values are $p = 0.384$ (in the high-cost treatment) and $p = 0.752$ (in the low-cost treatment). For the traded total buyer surplus, the p-values are $p = 0.198$ (in the high-cost treatment) and $p = 0.070$ (in the low-cost treatment).

Thus, none of the above variables differ significantly between periods 16-20 and 26-30 at conventional p-values, suggesting that our OO Markets converge to stable patterns already within 20 periods.

Appendix 7

Robustness of NO Market results with regard to add-on information

In the main treatments (HCT, CCT, SNT), the buyers in the NO Market saw only the overall net value (v_0) of all products in the market. In the case that buyers are only interested in the monetary aspects of the offered products, v_0 provides sufficient information for making a rational payoff-maximizing choice. But if the buyers were also interested in the add-on labels, the number of add-ons per se, or any other non-pecuniary information about the add-on features, their behavior might change if - in addition to information about v_0 – they are also given information about the offered products' add-ons.

In this context, it is important to recall the evidence on buyers' search and purchase behaviors in the OO Markets discussed in Section 3.2 and Appendix 3. This evidence indicates that buyers are neither interested in buying products with specific labels nor is their product demand affected by the number of add-ons per se. Moreover, the sellers anticipated these buyer behaviors. Based on these results, it seems unlikely that buyers will react to add-on labels or the number of add-ons per se in the NO markets. However, by withholding information about the products' add-ons in the NO Market of the main treatments, buyers' decision screen was less complex than their screen in the OO Market of the main treatments. One may speculate that this reduced complexity of the buyers' decision screen in the NO Market of the main treatments may have affected the adjustment to competitive outcomes in the NO Market. For this reason, when we ran the NO Markets in the robustness check treatments, we provided buyers with full information about the products' add-ons. Thus, the buyers' decision screen is similarly complex in the NO Market compared to the OO Market, with the only difference being that in the NO Market there is no longer any need to search.

Figure A.11 below shows buyers' surplus over time in the NO Markets of the main treatment of the Convex Cost Condition (CCT) and the robustness check treatment of the CCT. The figure clearly shows that both types of NO Markets smoothly converge close to the competitive outcome where buyers appropriate the whole surplus. Moreover, the adjustment process is very similar; it takes only a few periods for near-complete convergence. Thus, there is little reason to believe that variation in add-on information (and the associated variation in the complexity) of the buyers' screen had much impact on buyers' behavior in the NO Markets.

Figure A11: Average buyer surplus per trade as a share of the maximally possible surplus in NO Markets across the main and robustness treatments

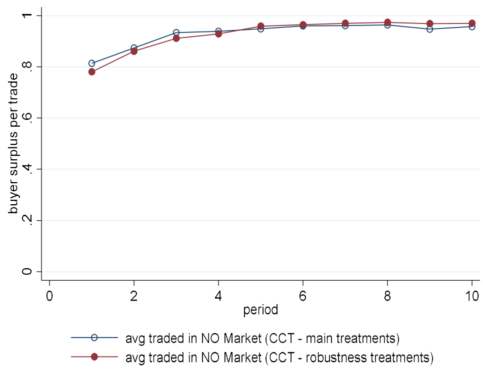

Appendix 8

The development of buyers' surplus in the OO Markets of the surplus-neutral treatment

Figure A12: Buyer surplus in OO Markets of the Surplus-Neutral Treatment

(a) High complexity markets

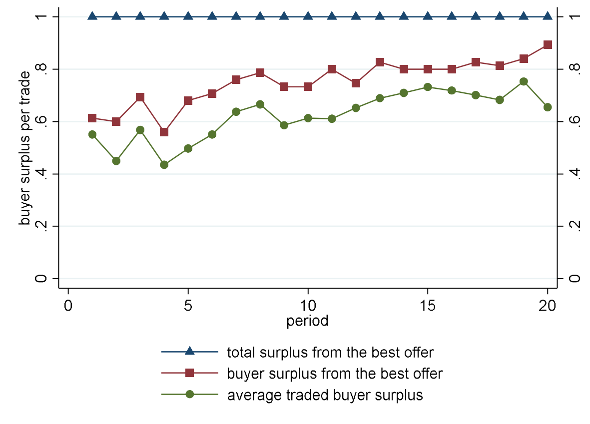

(b) Low complexity markets

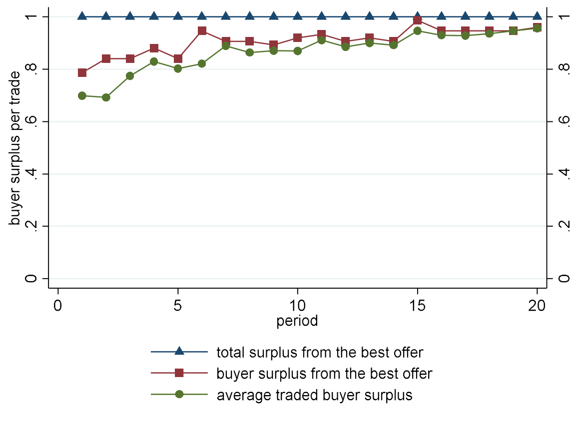

Notes: These figures show the average *buyer* surplus in traded products, the *buyer* surplus associated with the best available offer in the market, and the *total* surplus associated with the best offer in the OO Market. The figures are based on the data from the Surplus-Neutral Treatment. In this treatment, the total surplus from any (i.e., also the best) offer in the market is always 100% by construction because the number of chosen add-on features has no surplus consequences.

Figure A12b indicates that the buyer surplus converges to a very high level (94% of maximal total surplus) during periods 16-20 in low complexity markets. In contrast, Figure A12a shows that sellers in high complexity markets appropriate average profits of 30% of the maximal total surplus even in the long run (Period 16-20), which is only slightly lower than the sellers' surplus share of 32% in the treatments with surplus-enhancing extra features. Figure A12a also shows that there is a persistent gap in the buyers' surplus between the average traded and the best available product in the high complexity markets – a gap that is also significant during periods 16-20 (average gap is 12.8% with $p = 0.042$). This suggests that in the high complexity markets the buyers often fail to find the best available product. In contrast, Figure A12b indicates that this gap completely vanishes over time in the low complexity markets, i.e., after period 10 buyers almost always are able to identify and buy the best available product in this market.